

**Measurement of Charged Kaon Production by Neutrinos at
MINERvA**

by
Chris Marshall

Submitted in Partial Fulfillment
of the
Requirements for the Degree of
Doctor of Philosophy

Supervised by
Professor Kevin McFarland

Department of Physics and Astronomy
Arts, Sciences and Engineering
School of Arts and Sciences

University of Rochester
Rochester, New York

2016

Biographical Sketch

The author grew up in Plymouth, Minnesota. He attended Carleton College in Northfield, Minnesota, graduating *magna cum laude* in 2010 with a Bachelor of Arts degree in Physics. While an undergraduate, he began conducting particle physics research as an intern at the Soudan Underground Lab in the summer of 2008, and through the Science Undergraduate Laboratory Internship program at Fermilab in the summer of 2009.

The author began his graduate career at the University of Rochester in 2010 as a Ph.D. student, and joined the MINERvA collaboration under the supervision of Professor Kevin McFarland. He played an active role in a test beam experiment, detector calibrations, software development, and data analysis, culminating with the physics results described in this thesis.

Acknowledgements

It has truly been a privilege to work with the MINERvA collaboration and the Rochester neutrino group for the past six years. First and foremost, I would like to thank Kevin McFarland for many ideas, suggestions, questions, professional advice, and cups of espresso. I am grateful to have had outstanding mentorship from both Kevin and Steve Manly.

I thank Phil Rodrigues for his many contributions to this work, especially in the areas of statistics and making ROOT do what I want. I thank Jeremy Wolcott for his countless programming tips, and for his willingness to listen to me complain about the incompetence of the Minnesota Vikings. I thank Aaron Mislivec for many hours of helpful discussions on topics ranging from calibrations to covariance matrices.

It goes without saying in experimental particle physics that I could not have done it alone, and this work benefited greatly from many ideas and comments from MINERvA. I thank all of my collaborators over the years, especially Debbie Harris, Gabe Perdue, Dave Schmitz, and Hugh Gallagher. I thank Aiden Wang, who accomplished a summer's worth of work in a few weeks, then did that two more times and ended up paving the way to a really neat paper.

Finally, I thank Connie Jones, Laura Scuderi, and Laura Blumkin for their help with things too numerous to list.

Abstract

Charged kaon production by atmospheric neutrinos is a background in searches for the proton decay $p \rightarrow K^+ \bar{\nu}$. This channel is well motivated theoretically, and is favored by Supersymmetric Grand Unified Theories. Measurements of neutrino-induced K^+ production are important inputs for current and future proton decay searches at Super-K, Hyper-K and DUNE. The MINERvA neutrino-nucleus cross section experiment at Fermilab uses timing information to reconstruct K^+ decay-at-rest events. Differential cross section measurements are presented for both charged- and neutral-current neutrino interactions on plastic scintillator (CH). The cross section in K^+ kinetic energy, $d\sigma/dT_K$, is observed to be relatively flat between 0 and 500 MeV for both charged- and neutral-current scattering. Its shape is in good agreement with the prediction by the GENIE neutrino event generator when final-state interactions are included. A measurement of the hadronic visible energy in neutral-current K^+ events is also presented, as proton decay backgrounds come from events with very low visible energy. The spectrum is observed to be slightly softer in data than in GENIE. A search for charged-current coherent K^+ production, $\nu_\mu A \rightarrow \mu^- K^+ A$, is also presented. This rare process brings a K^+ on shell and leaves the target nucleus in its ground state. We find the first experimental evidence for the process at 3σ significance.

Contributors and Funding Sources

This work was supervised by a committee consisting of Advisor Kevin McFarland (Department of Physics), Professor Steve Manly (Department of Physics), Professor Gourab Ghoshal (Department of Physics), Professor David McCamant (Department of Chemistry), and chaired by Professor Kara Bren (Department of Chemistry).

From June 2015 through February 2016, this work was partially supported by the Fermilab URA visiting scholars program.

The MINERvA experiment is supported by the Fermi National Accelerator Laboratory under US Department of Energy contract No. DE-AC02-07CH11359 which included the MINERvA construction project. Construction support was also granted by the United States National Science Foundation under Award PHY-0619727 and by the University of Rochester. Support for participating scientists was provided by NSF and DOE (USA), by CAPES and CNPq (Brazil), by CoNaCyT (Mexico), by CONICYT (Chile), by CONCYTEC, DGI-PUCP and IDI/IGI-UNI (Peru), and by Latin American Center for Physics (CLAF). The MINERvA collaboration thanks the MINOS Collaboration for use of its near detector data, as well as the Fermilab staff responsible for the operation and maintenance of the NuMI beamline, MINERvA and MINOS detectors, and the physical and software environments that support scientific computing at Fermilab.

This thesis contains figures and text written by the author and previously published by the MINERvA collaboration. Chapters 6, 7, and 8 contain material from Refs. [1, 3, 2], and Chapter 4 contains material from Ref. [4]. A list of collaborators is found in Appendix B. Chapters 4 and 5 describe the MINERvA detector and simulation, and include descriptions of work done by other collaborators. The author's personal contributions to the calibration of the MINERvA detector are described in Sections 4.3, 4.4, and 4.6. The author was also responsible for maintaining the software and database infrastructure for detector calibrations. The author implemented the model described in Section 5.1.2, and the corrections to the kaon interaction model described in Section 5.2.1. The kaon-specific reconstruction in Section 6.2, as well as all of the data analysis described thereafter in Chapters 6, 7, and 8 was the work of the author.

Table of Contents

Biographical Sketch	ii
Acknowledgements	iii
Abstract	iv
Contributors and Funding Sources	v
List of Tables	xii
List of Figures	xiv
1 Why you should care if neutrinos produce kaons	1
2 Introduction	6
2.1 The Standard Model	13
2.1.1 The weak interaction	15
2.1.2 The strong interaction and QCD	18
2.2 Neutrinos in the Standard Model	22
2.3 Beyond the Standard Model	26
2.3.1 Grand Unification	26

2.3.2	Supersymmetry	29
2.4	Proton decay	30
2.4.1	Theory	30
2.4.2	Experimental searches	32
2.5	Neutrino interactions	40
2.6	Previous measurements of kaon production	47
2.7	Constraining proton decay backgrounds with MINERvA	49
3	Neutrino Beam and Flux	52
3.1	The NuMI beamline	52
3.1.1	Proton acceleration	52
3.1.2	The NuMI target	54
3.1.3	Focusing system	54
3.2	The neutrino flux prediction	57
3.2.1	The NA49 data	58
3.2.2	Extending to other interactions	59
3.2.3	Flux and uncertainties	60
3.2.4	Neutrino-electron scattering constraint	63
4	The MINERvA Detector	67
4.1	Readout and Electronics	70
4.2	Low-level calibrations	72
4.2.1	Pedestal calibration	72
4.2.2	FEB calibration	74
4.2.3	PMT gain calibration	74
4.2.4	Attenuation correction	75

4.3	Time slicing	76
4.4	Strip-to-strip calibration	78
4.4.1	Plane alignment correction	81
4.4.2	Plex and dead channels	84
4.4.3	Strip-level correction with the truncated mean	85
4.4.4	Plane-to-plane correction	87
4.5	Absolute energy calibration	90
4.6	Timing calibration	97
4.6.1	The reference track time	98
4.6.2	Electronics offsets	99
4.6.3	Time slewing correction	100
5	Simulation	103
5.1	Neutrino interaction simulation	103
5.1.1	Simulation of kaon production	105
5.1.2	Single kaon production in GENIE	107
5.2	Detector simulation	110
5.2.1	Geant4 simulation	111
5.2.2	Data overlay	114
5.2.3	Readout simulation	115
5.2.4	Optical model	116
5.2.5	PMT and ADC model	117
5.2.6	Timing simulation	119
6	Reconstruction and event selection	121
6.1	General reconstruction	122

6.2	Charged kaon reconstruction	125
6.2.1	Kinked track reconstruction	127
6.2.2	Time slivers	129
6.2.3	Low-energy K^+ scan	132
6.2.4	Additional K^+ cuts	135
6.3	Charged-current analysis	138
6.3.1	CC-specific selection cuts	139
6.3.2	Background subtraction	144
6.4	Unfolding	148
6.5	Efficiency correction	148
6.5.1	Systematic uncertainties	150
6.6	Neutral-current analysis	154
6.6.1	NC-specific selection cuts	155
6.6.2	Background subtraction	155
6.6.3	Unfolding	159
6.6.4	Efficiency correction	160
6.6.5	Systematic uncertainties	161
7	Results	165
7.1	Charged-current K^+ production	165
7.2	Neutral-current K^+ production	169
7.2.1	π^+ content of NC K^+ events	173
7.2.2	π^0 content of NC K^+ events	177
7.2.3	Toward a $p \rightarrow K^+\nu$ background prediction	179

8	Charged-current coherent K^+ production	182
8.1	Coherent K^+ theory	183
8.2	Event selection	187
8.3	Results	192
8.3.1	Cut-based analysis	193
8.3.2	Likelihood analysis	195
9	Conclusion	202
	Bibliography	204
	Appendices	213
A	Covariance matrices	214
B	The MINERvA Collaboration	217

List of Tables

2.1	List of important particles	15
2.2	Cherenkov threshold in water for different particle species.	38
3.1	FHC flux per POT	64
5.1	K^+ production in GENIE	106
6.1	Time slices and slivers	130
6.2	Summary of CC selection cuts	146
6.3	Signal and high non-kaon hadronic energy sideband	147
6.4	CC systematic uncertainties	153
6.5	CC knob twists	154
6.6	NC selection cuts	157
6.7	NC signal and sideband breakdown	159
6.8	NC $d\sigma/dT_K$ systematics	163
6.9	NC $d\sigma/dE_{vis}$ systematics	164
7.1	CC $d\sigma/dT_K$	168
7.2	Neutral-current $d\sigma/dT_K$ with absolute uncertainties	170
7.3	Neutral-current $d\sigma/dE_{vis}$ with absolute uncertainties	170

7.4	Michel electron rates in NC events	176
8.1	Coherent background event counts	190
9.1	Summary of agreement with GENIE	203
A.1	CC $d\sigma/dT_K$ flux covariance	214
A.2	CC $d\sigma/dT_K$ non-flux covariance	214
A.3	CC $d\sigma/dT_K$ statistical covariance	215
A.4	NC $d\sigma/dT_K$ flux covariance	215
A.5	NC $d\sigma/dT_K$ non-flux systematics covariance	215
A.6	NC $d\sigma/dT_K$ statistical covariance	215
A.7	NC $d\sigma/dE_{vis}$ flux covariance	216
A.8	NC $d\sigma/dE_{vis}$ non-flux systematics covariance	216
A.9	NC $d\sigma/dE_{vis}$ statistical covariance	216

List of Figures

2.1	Proton decay cartoons	10
2.2	Kaon FSI cartoons	12
2.3	Particles of the Standard Model	14
2.4	Weak interaction examples	16
2.5	$K_L^0 \rightarrow \mu^+ \mu^-$ suppression by the GIM mechanism	19
2.6	Hadronization diagram	21
2.7	Neutrino mass ordering	25
2.8	Charge screening due to vacuum polarization	27
2.9	Running coupling constants	30
2.10	Proton decay diagrams	31
2.11	Summary of proton decay limits	32
2.12	Super-Kamiokande ring patterns	34
2.13	Super-Kamiokande $K^+ \rightarrow \mu^+ \nu_\mu$ search	36
2.14	Super-Kamiokande $K^+ \rightarrow \pi^+ \pi^0$ search	37
2.15	Neutrino-electron scattering	41
2.16	Resonance pion production	42
2.17	Kaon production diagrams	44
2.18	Correlated nucleons in the nucleus	46

2.19	$K_S^0\Lambda$ production in a bubble chamber	48
3.1	The Fermilab accelerator complex	53
3.2	The NuMI target and cooling system	55
3.3	The NuMI focusing system	56
3.4	Weights from the NA49 experiment	59
3.5	The corrected neutrino flux prediction	62
3.6	Uncertainties on the flux prediction	63
3.7	Neutrino-electron flux tuning	65
3.8	Neutrino-electron flux constraint	66
4.1	The MINERvA detector	68
4.2	Scintillator planes	69
4.3	Scintillator strips	69
4.4	Pedestal ADC counts	73
4.6	Optical fiber routing	76
4.7	Time slicing	78
4.8	Rock muon event display	80
4.9	Relative energy variation in strips	81
4.10	A hypothetical rock muon intersects two adjacent strips.	82
4.11	Plane alignment fit with rock muons	83
4.12	Fraction of tracks with zero energy deposit	85
4.13	Iterative truncated mean energy	87
4.14	Truncated mean energy distribution	90
4.15	Peak energy by plane	91
4.16	Light level vs. time	92

4.17	Calibrated energy in data and simulation	93
4.18	True simulated cluster energy	94
4.19	Reconstructed cluster energy fit	95
4.20	Reconstructed vs. true energy	96
4.21	Time slewing corrections	98
4.22	Electronics timing offsets	100
4.23	Time slewing vs. pulse height	102
5.1	K^+ production in GENIE	108
5.2	Single kaon diagrams	109
5.3	Single kaon and associated production cross sections	111
5.4	K^+ -carbon scattering cross sections	112
5.5	Reweighted Geant4 K^+ -carbon cross section	115
5.6	PMT fiber weave and cross talk	118
5.7	Timing residual as a function of number of photoelectrons	120
6.1	Cluster classifications	123
6.2	MINOS match position resolutions	125
6.3	MINOS match timing resolution	126
6.4	Kaon event display example	127
6.5	Timing fit PDFs	129
6.6	Kinked track time gap	130
6.7	Time sliver event example	132
6.8	Time sliver gaps	133
6.9	Vertex time sliver event	134
6.10	K^+ energy vs. scan range	135

6.11 K^+ kinetic energy residuals	136
6.12 K^+ decay product energy	137
6.13 Cartoon of Michel background	138
6.14 Decay product number of hits	139
6.15 Cartoon of neutron background	140
6.16 Distance to decay product	140
6.17 Cartoon of pile-up background	141
6.18 True muon kinetic energy	142
6.19 Longest track range	143
6.20 MINOS charge reconstruction	144
6.21 Non-kaon hadronic energy	145
6.22 CC event distributions	147
6.23 CC T_K smearing matrix	149
6.24 Track and vertex efficiencies	150
6.25 Efficiency and $\Delta S = 1$	151
6.26 NC muon range	156
6.27 NC sideband shape comparisons	158
6.28 NC smearing matrices	160
6.29 NC efficiencies	161
6.30 Pie chart of background composition	163
7.1 Charged-current $d\sigma/dT_K$	166
7.2 GiBUU $d\sigma/dT_K$	167
7.3 Neutral-current $d\sigma/dT_K$	169
7.4 Neutral-current $d\sigma/dE_{vis}$	171

7.5	Water Cherenkov invisibility	173
7.6	K^+ energy spectrum at low E_{vis}	174
7.7	Michel electrons from π^+	175
7.8	NC K^+ with π^0	177
7.9	NC K^+ without π^0	178
7.10	MINERvA and atmospheric flux comparison	180
7.11	GENIE NC K^+ “invisible” cross section	181
8.1	Coherent K^+ Feynman diagram	184
8.2	Coherent K^+ and π^+ shape comparisons	186
8.3	Coherent $ t $ residual	189
8.4	Vertex energy and $ t $ distributions	190
8.5	Coherent signal and background event displays	191
8.6	Background and scan scale factor distribution	194
8.7	Expected background distribution	195
8.8	Coherent signal and background templates	196
8.9	$ t $ distribution and log-likelihood	197
8.10	Log-likelihood ratio to null hypothesis	198
8.11	Reconstructed coherent T_K	199
8.12	True $ t $ and K^+ energy for coherent events	200

Chapter 1

Why you should care if neutrinos produce kaons

We believe we have an understanding that all matter, the air we breathe, the coffee we drink, the monitor or piece of paper this is printed on, is composed of a set of elementary particles whose interactions we have studied and can describe. The goal of theoretical particle physics is to come up with models – basically sets of equations – that describe the behavior of those particles, and match up with what we observe in laboratory experiments and in the cosmos. The ultimate goal in this field would be to come up with a single theory – a single equation – that describes the whole universe. This “Theory of Everything” would have to have it all. It would have to explain the behavior of all the particles that we know about today. It would have to explain why some particles interact very strongly, and why others barely interact at all. It would have to explain the mysterious dark matter that astronomers see clustered around galaxies.

One type of theory, called a “Supersymmetric Grand Unified Theory” or SUSY

GUT, has many of these pieces. For over 40 years, theoretical physicists have been working on this kind of model. So far, experiments have not confirmed that a SUSY GUT is reality, but they haven't ruled it out either. The way to test these theories is to search for an otherwise forbidden process that they predict *must* occur, such as proton decay. Protons and their cousins, neutrons, form the nuclei of atoms, and are responsible for almost all of the mass that we can see in the universe. Neutron decay is common because the neutron is heavier than the proton, and can decay into a proton, an electron, and a neutrino. The proton doesn't have a lighter cousin, so if it were to decay it would have to be into a totally different kind of particle. Experiments have established that the proton lifetime is at least 10^{33} years, much longer than the age of the universe. SUSY GUTs say that a proton should preferentially decay into a positively-charged kaon and a neutrino.

Kaons are short-lived exotic particles, with a mass about half that of the proton. Protons and neutrons are each made up of three of the basic building blocks of the universe, called up and down quarks. Kaons are made of two quarks, and one of them is a different type, called a strange quark. A positively-charged kaon will decay in 12 nanoseconds. If you take a photograph of your feet, the light from the camera flash takes about that long to travel to the ground, reflect off your shoes, and travel back.

Neutrinos have masses at least 500 million times smaller than a proton, and possibly as much as 200 billion times smaller. They have no electric charge, and interact extremely weakly. On Earth, about 65 billion neutrinos from the sun pass through an area about the size of a penny every second. If you wanted to stop those neutrinos, you would need about a light-year of solid lead.

Observing a decay as rare as that of the proton is difficult. If you only had

one proton, you would have to watch it for 10^{33} years or more. The alternative employed by experiments is to watch 10^{34} protons. The current best experiment, called Super-Kamiokande in Japan, uses a tank with about 50 million kilograms of water. If one of the roughly 10^{34} protons in the tank were to decay in this way, a kaon and neutrino would suddenly appear in the middle of the tank. The neutrino would fly right out of the detector without interacting, and pass through the Earth and into space, never to be heard from again. The kaon would travel through the water, gradually losing energy until it comes to rest, and, roughly 12 nanoseconds later, decays. The kaon travels too slowly to be detected by Super-Kamiokande, but the products of its decay can be seen. It is as if the kaon escapes the scene, but leaves footprints. So all Super-Kamiokande has to do is wait, and see if they see the right footprints.

Unfortunately, proton decay is not the only way for those footprints to get into the detector. Particles whizzing around in space can bump into the Earth's atmosphere. These collisions produce spews of unstable particles, which decay into "atmospheric" neutrinos. Even though Super-Kamiokande is one kilometer underground, atmospheric neutrinos can go through the Earth and into the detector. Most of them, of course, will go right through the detector. But there are so many atmospheric neutrinos that occasionally one will hit an oxygen nucleus inside the water tank. When this happens, it is possible that a positively-charged kaon will be produced. If an atmospheric neutrino makes a positive kaon, and none of the other particles produced travel fast enough to be detected by Super-Kamiokande, then it looks exactly the same as what is expected from a proton decay. Kaon production by neutrinos is a "background" for proton decay.

The probability for a neutrino interaction to produce a kaon is small, but even-

tually such an event will occur, and it will look as if a proton has decayed. In order to control for this, we need to know how many atmospheric neutrinos pass through the detector, and how likely it is for one of them to scatter and make a kaon and nothing else. We already have pretty good measurements of the flux of atmospheric neutrinos, but it is much harder to measure the probability for a neutrino to produce a kaon, the “cross section” for kaon production. Right now, the estimate for how many events that look like proton decay but are really from neutrinos comes from a theoretical model. It is important to make sure that this model is correct by making a measurement of kaon production.

The MINERvA detector at Fermilab is made of solid plastic. When charged particles, including kaons, pass through MINERvA, they leave a trail of energy deposits. Based on the energy and time of these deposits, MINERvA is able to identify a charged kaon propagating in the detector, eventually stopping, and decaying at rest. MINERvA gets its neutrinos from the most intense neutrino beam in the world, and sees about one million times more kaons produced per year than what you would get from atmospheric neutrinos.

These results are the first high-statistics measurements of kaon production by neutrinos, with a factor of 50 more events than what previous experiments had observed. With the large sample of neutrino interactions, MINERvA has measured cross sections for different ways in which a neutrino can produce a kaon. Those cross section measurements are compared to predictions from theory, and found to agree surprisingly well. This is excellent news for proton decay experiments, which now have a benchmark for their background prediction.

As a by-product of these studies, we realized we could also search for an even rarer form of kaon production, where the neutrino scatters “coherently” off a carbon

nucleus and the nucleus remains intact instead of breaking apart. In our search, there are six events that look like coherent scattering, and the probability that it is just a statistical fluctuation of kaons that are produced by other processes is 0.3%. This type of neutrino interaction is predicted to occur, but prior to MINERvA no experiment had ever detected it. The observation of coherent kaon production validates our theoretical understanding of the coherent reaction.

Chapter 2

Introduction

The Standard Model of particle physics is a successful description of 17 known fundamental particles and their interactions. It explains the electromagnetic interaction that gives rise to electricity and magnetism. It explains the strong nuclear force that binds atomic nuclei together and overpowers the electrostatic repulsion of positively-charged protons. It explains the weak interaction observed in radioactive decays. Within the framework of the Standard Model, many predictions can be made and tested, in some cases with remarkable precision. For example, the prediction for the mass of the Z boson is 91.1874 ± 0.0021 GeV, and the current global best fit to experimental measurements is 91.1876 ± 0.0021 GeV [5].

However, we know the Standard Model is incomplete, and there is good reason to suspect that there is some greater theory that encompasses it. There are experimental observables that cannot be explained by the Standard Model, including gravity and dark matter. We know that roughly a quarter of the mass in the universe is due to something other than those 17 particles.

There are (at least) 26 parameters not predicted by theory in the Standard

Model, which must be determined experimentally in order to make predictions. Among those free parameters are particle masses, which inexplicably differ by 12 orders of magnitude. There are also free parameters that control the strength of the three interactions, which differ by about 6 orders of magnitude, also without explanation. (At least) 7 of those 26 parameters were added to account for massive neutrinos, which were not part of the original theory. In the original Standard Model, neutrinos were massless, and therefore could not oscillate from one flavor to another.

One class of proposed extensions to the Standard Model is called Grand Unification Theories (GUTs) [6, 7, 8, 9, 10]. The premise is that there exists some very large energy scale at which the “coupling constants” (and thus the strength) of the strong, electromagnetic, and weak interactions are equal. This larger symmetry breaks down at lower energies into the three distinct forces we observe. If Grand Unification were reality, it could explain the large difference in the strength of the forces, and perhaps also explain the smallness of the neutrino masses. New particles would have masses far too large to be created by accelerators, but Grand Unified Theories do have one important observable consequence: proton decay.

Another popular extension is Supersymmetry (SUSY), where each particle in the Standard Model is paired to a “superpartner.” This could help explain the puzzling “hierarchy” of masses, and the lightest superpartner is an excellent candidate for dark matter. It is popular [11, 12, 13, 14, 15, 16, 17, 18] to combine the two extensions into a SUSY GUT, a Supersymmetric Grand Unified Theory. Such a theory has nearly everything: a dark matter candidate, and explanation for the hierarchy, unification of the three distinct forces, and an explanation for neutrino masses.

The simplest GUTs (without SUSY) typically predict proton decay to a neutral pion and a positron, $p \rightarrow e^+\pi^0$, with a lifetime of order $10^{30} - 10^{31}$ years. The lifetime of this process is much longer (by about a factor of 10^4) in SUSY GUTs, but a faster process is allowed where a proton decays to a charged kaon and antineutrino, $p \rightarrow K^+\nu$, with typical lifetime predictions of order $10^{34} - 10^{35}$ years. Because of the compelling power of these theories to explain so many observables, experimentalists have been searching for proton decay since the early 1980s. To observe such a decay, one must either watch one proton for $\sim 10^{30}$ years, or 10^{30} protons for one year. There are about 10^{30} protons in 3,000 liters of water, or three tons of steel.

It began with simple experiments like Soudan I, a 30 ton iron tracking calorimeter in an abandoned iron mine, and KaimiokaNDE (the Kamioka Nucleon Decay Experiment), a 3,000 ton tank of water in a zinc mine. Charged particles traveling faster than the speed of light in water emit Cherenkov light [19], which is detected by light sensors lining the tank. When these experiments failed to observe proton decay, the next generations of detectors kept getting bigger and bigger. The current leader is Super-Kamiokande [20], a 50,000 ton water Cherenkov detector in the same zinc mine where KamiokaNDE once took data. It has ruled out $p \rightarrow e^+\pi^0$ with a lifetime shorter than 1.7×10^{34} years [21], invalidating many GUTs. Because of the larger kaon mass and the invisibility of the neutrino to experiments, it is harder to search for $p \rightarrow K^+\nu$ in a water Cherenkov detector. The current limit is 5.9×10^{33} years [22]. Furthermore, models that include SUSY predict longer lifetimes, many of which are still consistent with null experimental results. This has led to continued interest in $p \rightarrow K^+\nu$, and SUSY GUT models. In order to reach the 10^{34} year range, yet another increase in detector size is required. The proposed Hyper-Kamiokande [23] will have a total mass between 300,000 and 1 million tons

of water.

In a few years of running, Hyper-K will have an exposure of a few megaton-years, roughly the equivalent of watching 10^{36} protons for one year. If one of these protons, bound in an oxygen-16 nucleus, decays into a K^+ and $\bar{\nu}$, the residual excited-state nitrogen-15 nucleus promptly decays into its ground state, emitting a 6 MeV photon in the process. The $\bar{\nu}$ exits undetected. The K^+ is also not observed because its momentum is only ~ 339 MeV/c, and its velocity is slower than the speed of light in water. When that K^+ stops, about 2/3 of the time it decays into a μ^+ and a neutrino, each with a momentum of 236 MeV/c, with a lifetime of 12 ns. For a μ^+ , that momentum is above Cherenkov threshold, and detectable light is emitted. The μ^+ then decays $\sim 2 \mu\text{s}$ later into an electron and two unobserved neutrinos. The signal process is a 6 MeV photon, followed ~ 10 ns later by a μ^+ with precisely 236 MeV/c of momentum, followed $\sim 2 \mu\text{s}$ later by an electron.

The Hyper-K detector is shielded from cosmic rays by a kilometer of rock. Neutrinos produced in the atmosphere, however, will still reach the detector. Neutrinos have no electric charge and therefore do not produce Cherenkov light themselves, but the products of their interactions can mimic the signal of a proton decay. In particular, neutral-current interactions, where there is a neutrino in the final state, can produce K^+ with no other particles above threshold (as opposed to charged-current interactions, where the neutrino becomes a muon or electron, which can be detected). For example, a process like $\nu(^{16}\text{O}) \rightarrow \nu(^{15}\text{N})K^+\Lambda$ is often indistinguishable from $p \rightarrow K^+\nu$. The Λ baryon is an unstable particle that decays in ~ 0.3 ns to either $p\pi^-$ or $n\pi^0$. When it decays $\Lambda \rightarrow n\pi^0$, the π^0 decays in about 1 fs to two photons which are easily detected by Hyper-K. The decay to $p\pi^-$ however, is undetectable when π^- momentum is below 200 MeV/c. In that case, there is




Figure 2.1: A cartoon of a proton decay to a K^+ and neutrino in a large detector (left), and a potential irreducible background event from an atmospheric neutrino (right). Neutral particles are drawn with dashed lines, and the background event is identical to the proton decay if the proton and pion from lambda decay do not produce Cherenkov light.

no additional particle to make it obvious that the K^+ came from an atmospheric neutrino rather than a proton decay. You have the same nitrogen-15 nucleus and K^+ that you are looking for. An example proton decay and neutrino-induced event are shown in Figure 2.1.

So far, in the smaller Super-K, the total published exposure has been 260 kiloton-years, and no proton decay-like events have been observed [22]. The current best prediction for the expected number of neutrino-induced background events is 0.25. That number is determined by taking a prediction for the flux of atmospheric neutrinos and antineutrinos [24], and a theoretical model of the neutrino K^+ production cross section [25], the probability for a neutrino to interact with oxygen and make a K^+ . If the cross section model is wrong, the expected number of background events could be much smaller, or much larger than 0.25.

For a Hyper-K exposure of several megaton-years, that background prediction rises to a few events. If 10 are observed, have you discovered proton decay? What

is the uncertainty on that prediction of a few events? If the actual number expected is 3, then the probability of observing 10 or more is 1 in 1,000. However, if the cross section model is wrong by a factor of 2 and the number expected is 6, then there is an 8% chance of observing 10 or more. This difference is very important for the next generation of proton decay searches.

A search for $p \rightarrow K^+\nu$ will also be performed at DUNE [26], a liquid argon time projection chamber (LAr TPC) scheduled to start taking data in 2025. Unlike in a water Cherenkov detector, a LAr TPC can detect the K^+ from proton decay and measure its momentum. It can also detect other low-energy hadrons that are produced in neutrino reactions that would be absent from $p \rightarrow K^+\nu$, such as the products of $\Lambda \rightarrow p\pi^-$ decay. DUNE is expected to have similar sensitivity to $p \rightarrow K^+\nu$ as Hyper-K despite a much smaller mass of 40 kilotons. This sensitivity is achieved by a very high efficiency to identify a $p \rightarrow K^+\nu$ event, and an extremely low background rate from atmospheric neutrinos.

In DUNE, the K^+ momentum can be measured, and compared to the expected momentum for $p \rightarrow K^+\nu$. For a free proton at rest, it is a simple two-body decay and the K^+ momentum is always 105 MeV/ c . For a proton bound in argon, however, the momentum spectrum is smeared by Fermi motion, the momentum distribution of protons inside the nucleus. This means that the K^+ momentum could be anywhere from a few 10s to a few 100s of MeV/ c . The spectrum is further distorted by “final-state interactions” (FSI). In bound proton decay, the K^+ is born inside the nucleus, and can undergo interactions with other nucleons on its way out. This would reduce the K^+ momentum observed in the detector, and could result in nucleons being ejected from the nucleus. This is illustrated in Figure 2.2. Modeling FSI is important for both signal and background predictions for $p \rightarrow K^+\nu$ in DUNE.




Figure 2.2: For a bound proton decay, the K^+ is born inside the nucleus, with momentum smeared by the motion of the initial proton (left). On its way out of the nucleus, it can interact with other nucleons (right).

MINERvA is a fine-grained detector that records the energy deposited by all charged particles as they move through the solid hydrocarbon plastic scintillator. MINERvA also measures the time of these deposits with a resolution of 3 ns, making it possible to identify the products of a K^+ decay-at-rest based on timing. This reconstruction is used to select samples of both charged-current and neutral-current neutrino-induced K^+ events. With these samples, the rate of neutral-current K^+ production is measured as a function of the K^+ energy, and as a function of the total detected energy of other particles. The charged-current sample is used to measure the K^+ energy spectrum, which is sensitive to the same kaon-nucleus final-state interactions that would affect K^+ from proton decays. Together with a measurement of the atmospheric neutrino flux, these neutrino cross section measurements constrain background predictions in searches for $p \rightarrow K^+\nu$, and constrain models of FSI that affect the K^+ spectrum.

In this chapter, the Standard Model is discussed in greater detail in Section 2.1, and the role of neutrinos is discussed in Section 2.2. Section 2.3 describes the short-

comings of the standard model and their potential solutions, namely Supersymmetry and Grand Unification. Proton decay models are introduced in Section 2.4. Neutrino interactions are discussed in Section 2.5, with special attention to processes which mimic the experimental signature of the proton decay $p \rightarrow K^+\nu$. Previous measurements of K^+ production by neutrinos are described in Section 2.6, and MINERvA's role is introduced in Section 2.7.

2.1 The Standard Model

The particles of the Standard Model (SM) are divided into two categories, fermions and bosons, based on their intrinsic spin. The fundamental bosons, which have integer spin, are also called “gauge bosons” or “force carriers” because they are responsible for mediating the strong (gluon), weak (W^\pm and Z), and electromagnetic (photon) interactions.

The fermions, which have half-integer spin, are further subdivided based on the strength of their interactions by the three SM forces. Quarks are fermions which couple to gluons, and therefore feel the strong force. They are split into “up-type” quarks (up, charm, top), which have electric charge $2/3$, and “down-type” quarks (down, strange, bottom), which have charge $-1/3$. Leptons are fermions which do not interact via the strong force. They are further divided into the charged leptons (electron, muon, tau), which have charge -1 , and the neutral neutrinos. In addition to these divisions based on interactions, SM fermions are also divided into three “families.” Mass is the only difference between fermions in the first, second, and third family of their respective type. The particles of the SM are shown in Figure 2.3. Each particle has a corresponding antiparticle with opposite charge.

mass →	2.4 MeV/c ²	1.27 GeV/c ²	171.2 GeV/c ²	0	≈126 GeV/c ²
charge →	2/3	2/3	2/3	0	0
spin →	1/2	1/2	1/2	1	0
	u up	c charm	t top	γ photon	H Higgs boson
QUARKS	4.8 MeV/c ² -1/3 1/2 d down	104 MeV/c ² -1/3 1/2 s strange	4.2 GeV/c ² -1/3 1/2 b bottom	0 0 1 g gluon	
	0.511 MeV/c ² -1 1/2 e electron	105.7 MeV/c ² -1 1/2 μ muon	1.777 GeV/c ² -1 1/2 τ tau	91.2 GeV/c ² 0 1 Z Z boson	
LEPTONS	<2.2 eV/c ² 0 1/2 ν_e electron neutrino	<0.17 MeV/c ² 0 1/2 ν_μ muon neutrino	<15.5 MeV/c ² 0 1/2 ν_τ tau neutrino	80.4 GeV/c ² ±1 1 W W boson	GAUGE BOSONS

Figure 2.3: The 17 particles of the standard model subdivided into categories: bosons, fermions (quarks, leptons).

The quarks are not observed individually, a phenomenon discussed further in Section 2.1.2. Quarks are commonly observed in quark-antiquark bound states called “mesons,” and three-quark bound states called “baryons.” There is experimental evidence for the existence of four-quark states [27, 28], and possibly even five-quark states [29]. However, these exotic states are rare, difficult to produce, and short-lived, and are not relevant for neutrino scattering. Important particles for this thesis are given in Table 2.1 with masses, lifetimes, and decay modes.

Particle	Quark composition	Mass (MeV/ c^2)	Lifetime (ns)	Decay modes (%)
μ^-		105.7	2197.0	$e^- \bar{\nu}_e \nu_\mu$ (100)
π^+	$u\bar{d}$	139.6	26.0	$\mu^+ \nu_\mu$ (99.99)
π^0	$\frac{u\bar{u}-d\bar{d}}{\sqrt{2}}$	135.0	10^{-7}	$\gamma\gamma$ (98.8)
K^+	$u\bar{s}$	493.7	12.4	$\bar{\nu}_\mu \mu^+$ (64) $\pi^+ \pi^0$ (21) $\pi^+ \pi^+ \pi^-$ (6) $e^+ \nu_e \pi^0$ (5)
K_S^0	$\approx \frac{d\bar{s}-s\bar{d}}{\sqrt{2}}$	497.6	0.1	$\pi^0 \pi^0$ (31) $\pi^+ \pi^-$ (69)
K_L^0	$\approx \frac{d\bar{s}+s\bar{d}}{\sqrt{2}}$	497.6	51.2	$\pi^\pm e^\mp \nu_e$ (41) $\pi^\pm e^\mp \nu_e$ (27) $\pi^0 \pi^0 \pi^0$ (20) $\pi^+ \pi^- \pi^0$ (12)
Λ	uds	1115.7	0.3	$p\pi^-$ (64) $n\pi^0$ (34)
Σ^+	uus	1189.4	0.1	$p\pi^0$ (52) $n\pi^+$ (48)
Σ^-	dds	1197.4	0.1	$n\pi^-$ (100)

Table 2.1: Masses, lifetimes, and decay modes for particles important to this thesis. Neutral kaons interact in states with definite strangeness ($d\bar{s}$ or $\bar{d}s$, but propagate and decay in the states shown in the table.

2.1.1 The weak interaction

The weak interaction is mediated by the exchange of massive gauge bosons, the W and Z . Unlike other force carriers, the W boson is charged, and gives rise to “charged-current” (CC) weak interactions. The W connects up-type and down-type quarks, which differ in charge by one unit. Neutrinos also interact with the W , turning the neutrino into its charged lepton partner or vice versa. Figure 2.4 shows two Feynman diagrams of weak interactions. The first is simple beta decay, $n \rightarrow e^- \bar{\nu}_e p$, a neutron decaying to a proton, and electron, and an electron antineutrino. The second is a simple neutrino scatter, in which a muon neutrino interacts with a down quark to produce an up quark and μ^- . When two “spectator” quarks are




Figure 2.4: Feynman diagrams of a beta decay (left), and a muon neutrino quasi-elastic interaction (right).

added, this process is $\nu_\mu n \rightarrow \mu^- p$, called charged-current quasi-elastic scattering.

The vertices in both of the diagrams in Figure 2.4 are essentially $d + W^+ \rightarrow u$ and $\ell + W^+ \rightarrow \nu_\ell^-$, where $\ell = e, \mu$. A W^+ turns a down-type quark into an up-type quark, or a (negatively) charged lepton into a neutrino. The eigenstates of the weak interaction are “doublets” of Standard Model particles:

$$\begin{pmatrix} u \\ d \end{pmatrix}, \begin{pmatrix} c \\ s \end{pmatrix}, \begin{pmatrix} t \\ b \end{pmatrix}, \begin{pmatrix} e \\ \nu_e \end{pmatrix}, \begin{pmatrix} \mu \\ \nu_\mu \end{pmatrix}, \begin{pmatrix} \tau \\ \nu_\tau \end{pmatrix}. \quad (2.1)$$

There is mixing between the quark mass eigenstates (which propagate through space) and the quark weak eigenstates (which interact via the weak force). This mixing is given by the Cabibbo-Kobayashi-Masakawa (CKM) matrix [30], named for Nicola Cabibbo, Makoto Kobayashi, and Toshihide Masakawa. The relationship between the mass states (d', s', b') and the weak states (d, s, b) is:

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix}, \quad (2.2)$$

where $|V_{ij}|^2$ is proportional to the probability for a quark of flavor i to decay into a quark of flavor j . Mixing to and from the third family is very small ($|V_{tb}| = 0.999146$ [5]), so it is convenient to express the mixing as a 2×2 matrix, which can then be parameterized in terms of the Cabibbo angle θ_C as

$$U_{CKM} = \begin{pmatrix} \cos \theta_C & \sin \theta_C \\ -\sin \theta_C & \cos \theta_C \end{pmatrix}. \quad (2.3)$$

The amplitude of “Cabibbo-favored” transitions such as $d + W^+ \rightarrow u$ or $\bar{u} + W^+ \rightarrow \bar{d}$ has a factor of $\cos \theta_C$, while the amplitude for “Cabibbo-suppressed” transitions such as $d + W^+ \rightarrow c$ or $\bar{u} + W^+ \rightarrow \bar{s}$ picks up a factor of $\sin \theta_C$. The Cabibbo angle θ_C is measured experimentally to be 13.02° [5]. The magnitude of the suppression is $\tan^2 \theta_C = 0.0535$. This means that Cabibbo-favored processes are about a factor of 20 more likely to occur than their Cabibbo-suppressed counterparts.

Weak “neutral-current” (NC) reactions also occur, and proceed via Z boson exchange. The weak neutral current was discovered by observing neutrino interactions $\nu_\mu N \rightarrow \nu_\mu X$ and $\nu_\mu e^- \rightarrow \nu_\mu e^-$ at the Gargamelle bubble chamber [31]. Charged particles can also undergo weak NC interactions, but the same processes can occur via the electromagnetic interaction, so at low energies (much smaller than the Z mass) the weak contribution is small. However, weak NC processes are very important in neutrino-nucleus scattering.

Weak neutral current reactions do not change the quark flavor, as flavor-changing

neutral currents are suppressed by the Glashow-Iliopoulos-Maiani (GIM) mechanism [32]. A famous example of this mechanism is shown in Figure 2.5. Without additional quark families and mixing, the neutral kaon decay $K_L^0 \rightarrow \mu^+ \mu^-$ should occur at a much higher rate than was observed in experiments. Prior to the discovery of the charm quark, GIM proposed its existence in order to explain the observed low rate. In the diagram with a u quark, the amplitude picks up a factor of $\sin \theta_C$ at the vertex labeled 1, and $\cos \theta_C$ at 2. In the diagram with a c quark, the factors are $\cos \theta_C$ and $-\sin \theta_C$ at vertices 1 and 2, respectively. The result is that the two amplitudes nearly cancel. The cancellation would be perfect except for the difference in the masses of the u and c quarks.

The weak interaction is the only one to violate strangeness. In strong interactions, the net number of strange quarks $S = N_{\bar{s}} - N_s$ is conserved. The K^+ is the lightest strange particle, and therefore can only decay weakly.

2.1.2 The strong interaction and QCD

The strong interaction is described by the theory of quantum chromodynamics (QCD). Quarks and gluons have a quantum number called “color charge,” the QCD analog of electric charge in the electromagnetic interaction. There are three “colors,” typically called red, green, and blue although color in this sense has nothing at all to do with reflecting light of some wavelength and the “colors” could just as easily be called strong charges one, two, and three. Unlike the electromagnetic interaction, in which the mediating photon is electrically neutral, gluons themselves have color charge. This leads to the phenomenon of color confinement, which means that single quarks or gluons cannot be isolated – the particles that we observe in nature are color-neutral.




Figure 2.5: The “box diagram” for the decay $K_L^0 \rightarrow \mu^+ \mu^-$ in the Standard Model. Two diagrams contribute, with either a u or c quark as indicated. The amplitudes cancel, except for a small difference due to the masses of the u and c quarks, suppressing the process.

There are two common and well established types of color-neutral particles composed of quarks (called “hadrons”). Mesons are made of one quark and one antiquark, while baryons are made of three quarks. For a meson to have no color, the color of the antiquark must cancel out the color of the quark. For example, a red up and an anti-red down can combine to form a colorless π^+ . For a baryon to have no color, it must have one quark of each color. For example, a proton is made of two up quarks and a down, and must have one red, one blue, and one green quark.

When a single quark is liberated from a hadron at high energies, say by a neutrino striking a down quark within a neutron, a virtual gluon field binds this quark

to its partners. Unlike the electromagnetic force, which weakens with distance, the strength of this field does not deteriorate. The struck quark and the two spectators want badly to be part of colorless objects, so a $q\bar{q}$ pair is “pulled from the vacuum” when a radiated gluon pair produces. The new quark can combine with the spectators to form a colorless baryon, while the new antiquark can combine with the struck quark to form a meson. In Figure 2.6, this is shown with a $d\bar{d}$ pair forming a π^+ and neutron. This process is called hadronization. At very high energies, such as those of particle colliders like the LHC, a quark will hadronize into many nearly collinear mesons which are observed in a detector as “jets.” In neutrino scattering at a few GeV, we may see one or more pions in the detector as a result of this hadronization.

Another interesting feature of QCD is asymptotic freedom, which makes the quark-gluon interaction strength asymptotically weak at high energies and short distances. As the coupling constant α_s is small, higher-order terms in α_s can be neglected, making it possible to calculate QCD processes using perturbation theory. The high energy regime where perturbation theory works well is of great interest in collider physics. At the LHC, for example, QCD is used to calculate backgrounds in searches for the Higgs boson and other phenomena beyond the standard model (see Section 2.3). These backgrounds arise from qq , qg , and gg interactions at high energies that proceed via known SM processes. When the energy is sufficiently high, these processes can be calculated to first or second order in α_s and such calculations give good agreement with data.

At lower energies, however, perturbative QCD breaks down. There are no analytic solutions to QCD at low energy. Lattice QCD, where quark fields are calculated on a lattice of points rather than a continuum, can be a good approximation if the lattice is sufficiently fine. However, a fine lattice is computationally expensive. In-




Figure 2.6: A neutrino interacts with a quark inside a neutron. The struck quark radiates a gluon, which produces a quark-antiquark pair. These new quarks combine to form colorless hadrons.

stead, it is necessary to use effective theories, which describe some specific regime of phase space adequately, or to use parameterizations of data.

One type of effective theory that is important in neutrino physics is the nucleon form factor. At low energies where α_s is large, the interaction of a photon, W , or Z boson with a nucleon is very complicated. Inside the nucleon, quarks can radiate gluons, which can fluctuate to quark-antiquark pairs called the quark sea. This is similar to the hadronization shown in Figure 2.6, but occurs inside the nucleon. Calculating all of the possible interactions with quarks, including sea quarks, and gluons, is impossible. Instead, a simple expression is used, and parameters are

extracted from data. This makes it possible to represent the data without modeling all of the details of the messy interaction.

2.2 Neutrinos in the Standard Model

The original SM contained only massless neutrinos. The neutrinos fit nicely into this model, with each of the three charged leptons combining with a neutrino to form a doublet as in Equation 2.1. Neutrinos interact via the W boson only with their partner; a ν_ℓ ($\bar{\nu}_\ell$) charged-current interaction will always yield a charged lepton ℓ^- (ℓ^+), and there is no mixing.

The observation of neutrino oscillations by Super-Kamiokande [33] and later the Sudbury Neutrino Observatory (SNO) [34, 35] proved this to be incomplete. The neutrino “flavor” states of Equation 2.1 do not have definite mass, and therefore do not have definite momentum at fixed energy. They are not eigenstates of the time evolution operator, and thus cannot propagate. Instead, neutrinos interact in states of distinct flavor, $|\nu_\alpha\rangle$ for $\alpha = e, \mu, \tau$ but propagate in states of distinct mass $|\nu_i\rangle$ for $i = 1, 2, 3$. They are related by a mixing matrix, U , similar to the CKM matrix, called the Pontecorvo-Maki-Nakagawa-Sakata matrix:

$$|\nu_\alpha\rangle = \sum_{i=1}^3 U_{\alpha i} |\nu_i\rangle. \quad (2.4)$$

The time evolution is given by

$$|\nu_\alpha(t)\rangle = \sum_{i=1}^3 U_{\alpha i} e^{-iE_i t} |\nu_i\rangle, \quad (2.5)$$

where E_i is the total energy of the neutrino. At time $t = 0$ Equation 2.5 trivially

collapses into Equation 2.4. The amplitude to observe the neutrino at a later time in a flavor eigenstate $|\nu_\beta\rangle$, where β may be but is not necessarily α , is

$$\begin{aligned}
 \langle \nu_\beta | \nu_\alpha(t) \rangle &= \langle \nu_\beta | \sum_{i=1}^3 U_{\alpha i} e^{-iE_i t} | \nu_i \rangle \\
 &= \left(\sum_{j=1}^3 U_{\beta j}^* \langle \nu_j | \right) \left(\sum_{i=1}^3 U_{\alpha i} e^{-iE_i t} | \nu_i \rangle \right) \\
 &= \sum_{i=1}^3 U_{\alpha i} U_{\beta i}^* e^{-iE_i t}
 \end{aligned} \tag{2.6}$$

where the last equality uses $\langle \nu_j | \nu_i \rangle = \delta_{ij}$. In an approximation where there are only two neutrinos, the mixing matrix U can be expressed in terms of a mixing angle θ , which is essentially a rotation between the mass and flavor bases:

$$U = \begin{pmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{pmatrix}. \tag{2.7}$$

Note the similarity to the two-flavor quark mixing parameterization of Equation 2.3. Equation 2.6 simplifies to

$$\begin{aligned}
 \langle \nu_\beta | \nu_\alpha(t) \rangle &= \langle \nu_\beta | (\cos \theta e^{-iE_1 t} | \nu_1 \rangle + \sin \theta e^{-iE_2 t} | \nu_2 \rangle) \\
 &= (-\sin \theta \langle \nu_1 | + \cos \theta \langle \nu_2 |) (\cos \theta e^{-iE_1 t} | \nu_1 \rangle + \sin \theta e^{-iE_2 t} | \nu_2 \rangle) \\
 &= -\sin \theta \cos \theta e^{-iE_1 t} + \sin \theta \cos \theta e^{-iE_2 t},
 \end{aligned} \tag{2.8}$$

and the oscillation probability obtained by squaring the amplitude is

$$\begin{aligned}
P(\nu_\alpha \rightarrow \nu_\beta) &= \sin^2 \theta \cos^2 \theta (2 - 2 \cos(E_2 - E_1)t) \\
&= \sin^2 2\theta \sin^2 \frac{(E_2 - E_1)t}{2}.
\end{aligned} \tag{2.9}$$

For ultra-relativistic neutrinos, $E \gg m$ and $t \sim L$, the distance traveled by the neutrinos. The three-momentum p is equal for all mass components of the ν_α state, and the energy E_i can be Taylor expanded as (in natural units with $\hbar = c = 1$)

$$\begin{aligned}
E_i &= \sqrt{p^2 + m_i^2} \\
&= p + \frac{m^2}{2p}.
\end{aligned} \tag{2.10}$$

For very light neutrinos $p \sim E$ and the oscillation probability in Equation 2.9 becomes

$$P(\nu_\alpha \rightarrow \nu_\beta) = \sin^2 2\theta \sin^2 \frac{\Delta m^2 L}{4E}, \tag{2.11}$$

where $\Delta m^2 = m_2^2 - m_1^2$, the difference of the squared masses of the two neutrinos. In Equation 2.11, the argument $\frac{\Delta m^2 L}{4E}$ is expressed in natural units. It is equal to $1.27 \frac{\Delta m^2 L}{E}$ when Δm^2 is expressed in eV^2 , and the ratio L/E is expressed in km/GeV or, equivalently, m/MeV . Two important consequences can be observed from Equation 2.11. First, neutrino oscillation implies that neutrinos are massive. If Δm^2 were identically zero, then the neutrinos would not oscillate, and for there to be a nonzero difference in the neutrino masses at least one of the states must have nonzero absolute mass. Second, the probability for a neutrino initially in one flavor state to be observed in another flavor state oscillates as a function of L/E , the ratio

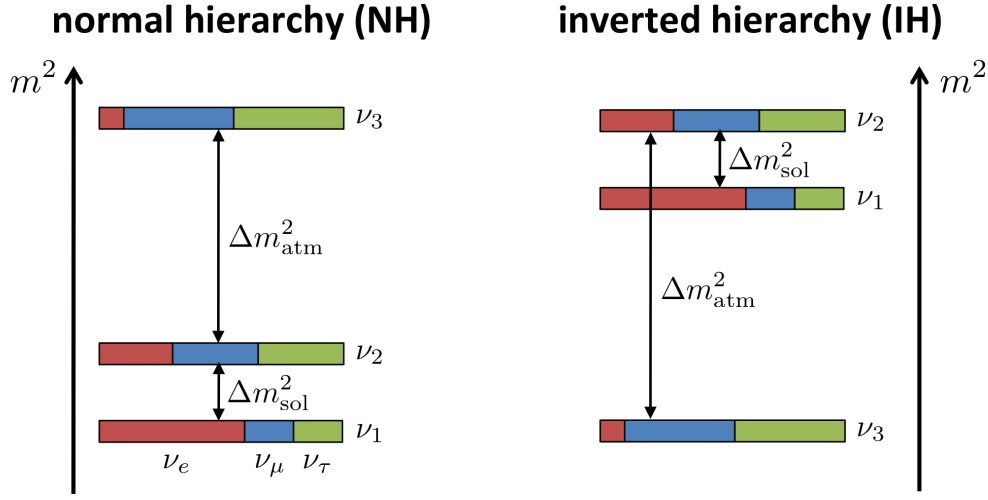


Figure 2.7: The neutrino mass splittings are shown for the normal ordering (left) and the inverted ordering (right), with the color composition of each bar corresponding to the flavor composition of each mass eigenstate. Taken from Ref. [36]

of the distance traveled to the energy of the neutrinos.

In the full three-flavor picture, there are two independent mass splittings. The smaller one, $\Delta m_{12}^2 = (7.54_{-0.22}^{+0.26}) \times 10^{-5} \text{ eV}^2$ [5], is often called the “solar” splitting because it is primarily responsible for the oscillations of $\sim 1 \text{ MeV}$ neutrinos from the sun. The larger splitting is relevant for the oscillation of $\sim 1 \text{ GeV}$ atmospheric neutrinos. It is known from matter effects that ν_2 is heavier than ν_1 , but vacuum oscillations are sensitive only to \sin^2 of the squared mass splitting, so the sign of Δm_{31}^2 is not known. Its absolute value is experimentally $2.4 \times 10^{-3} \text{ eV}^2$ [5]. The magnitudes of the oscillations are parameterized in terms of three independent mixing angles, called θ_{12} , θ_{23} and θ_{13} . A diagram of the mass splittings showing the flavor content of each of the mass eigenstates is shown in Figure 2.7.

2.3 Beyond the Standard Model

2.3.1 Grand Unification

The Standard Model lacks an explanation for why the strengths of the strong, weak, and electromagnetic forces are so different. Equivalently, it fails to explain why the coupling of fermions to gluons, W and Z bosons, and photons is so different. Grand Unification Theories (GUTs) provide such an explanation. While there are many such models [6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 17, 18], they all propose a scale at which a larger symmetry unifies the three forces with a single coupling constant. This scale is called the “GUT scale” and is typically $10^{15} - 10^{16}$ GeV. At energies much below the GUT scale, the larger symmetry breaks down into the three forces we observe. GUT models predict new bosons, with masses at the GUT scale. While these bosons are far too heavy to be produced in particle colliders, they may have played a role in the early universe.

Coupling constants are not really constants. Instead, the strength of the interactions described by these “constants” changes very slowly with energy. The electromagnetic interaction grows stronger with energy due to charge screening. The vacuum can fluctuate to a virtual e^+e^- pair, borrowing energy for the particle masses for a very short time governed by the time-energy relation, $\Delta E \Delta t \sim \hbar$. The vacuum can fluctuate to pairs of other particles, but the allowed time of the fluctuation is inversely proportional to the particle mass, so e^+e^- pairs dominate. A charged particle, for example an electron, will polarize the vacuum, attracting the opposite-charge virtual particles. These virtual particle pairs act like dipoles in the electric field of the charged particle. As a result, the observed charge of the electron varies as a function of distance. For a low-energy, long-wavelength probe,




Figure 2.8: Vacuum fluctuations to e^+e^- pairs act like dipoles, alining with field lines of a charged particle and screening the charge.

the electromagnetic interaction is weaker because some of the electron's charge is screened by the polarized vacuum. As the energy of the probe increases, the wavelength decreases, and the interaction strengthens. Figure 2.8 is an illustration of dipole-like e^+e^- pairs alining with the field lines of an electron, demonstrating the charge screening effect.

The color charge of quark-antiquark vacuum fluctuations gives rise to a similar screening effect for the strong coupling. However, there is a larger effect with the opposite sign due to the fact that the force carrying gluons have color charge themselves. This gives rise to the asymptotic freedom discussed in Section 2.1.2 and is also responsible for the strong coupling decreasing at high energies. The “beta

function” that describes how the strength of the strong interaction changes with energy is the slope as a function of the logarithm of the energy scale, and for QCD in an arbitrary number of dimensions and with any number of quark flavors is given by

$$\beta(\alpha_s) = \frac{\alpha_s^2}{\pi} \left(-\frac{11N}{6} - \frac{n_f}{3} \right), \quad (2.12)$$

where N is the number of dimensions, and n_f is the number of quark flavors. For three dimensions and six quark flavors, the slope is negative [37, 38].

Quarks and leptons are also unified at the GUT scale in this type of model. New bosons are introduced, typically called X and Y , with masses at the GUT scale. The bosons are fractionally charged. The X has a charge of $4/3$ and decays by $X \rightarrow uu$ and $X \rightarrow e^+\bar{d}$. The Y has a charge of $1/3$ and decays by $Y \rightarrow e^+\bar{u}$, $Y \rightarrow du$, and $Y \rightarrow \bar{d}\bar{\nu}_e$. In these decays, the quark and lepton families are interchangeable, so u , d , e , and ν_e could be instead the second or third family.

Interactions mediated by the X and Y bosons do not conserve baryon or lepton number, but conserve the difference $B - L$. Since baryon number is not conserved, the proton becomes unstable. Thankfully, the proton lifetime is proportional to the fourth power of the GUT scale, $\tau \propto M_{GUT}^4$. Since M_{GUT} is of order 10^{15} GeV, the proton lifetime is many orders of magnitude longer than the age of the universe. The other consequence of the large GUT scale is that the X and Y bosons are far beyond the reach of colliders like the LHC, or even hypothetical future colliders. Proton decay is the most straightforward testable prediction. Its observation would be very strong evidence for Grand Unification. The specific decay mode and lifetime can be used to discriminate between the many possible GUT models.

2.3.2 Supersymmetry

Supersymmetry (SUSY) is a class of extensions to the Standard Model that unify fermions and bosons. Each SM fermion has a bosonic “superpartner,” and vice versa. If SUSY were an exact symmetry, the masses of the superpartners would be identical to those of the SM particles. This cannot be the case, as superpartners as light as an electron would be easy to detect, and so far there is no evidence for these particles up to energies of hundreds of GeV. Therefore SUSY must be spontaneously broken so that the superpartner masses are in the allowed region.

SUSY could explain the “hierarchy problem,” why the weak force is $\sim 10^{32}$ times stronger than gravity. It could also provide a candidate for dark matter, as the lightest superpartner is stable in many SUSY models. The additional particles also increase the screening effect described in Section 2.3.1 for all three SM forces. This modifies the slopes of the couplings as a function of energy. One downside to “bare” GUTs is that the couplings for the three forces, when extrapolated to the GUT scale, do not meet at a point. Combining a GUT with SUSY, however, makes it possible to truly unify the three couplings. Due to the slope change, it also increases the energy at which the three couplings are unified by about an order of magnitude, as shown in Figure 2.9.

The superpartners are called “sparticles.” Bosonic superpartners of SM fermions are named by prepending an s to the name of the fermion. The squarks are *sup*, *sdown*, *scharm*, *sstrange*, *sbottom*, and *stop*. The sleptons are the *selectron*, *smuon*, and *stau*, with corresponding *sneutrinos*. Fermionic superpartners of SM bosons are diminutive in Italian. The force carriers are the *gluino*, *photino*, *Wino*, and *Zino*. The Higgs boson has a superpartner as well, the *Higgsino*.

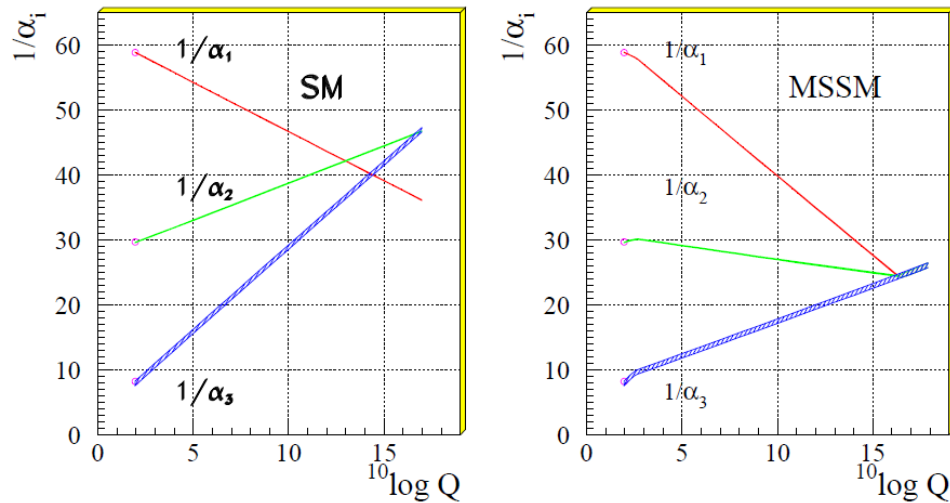


Figure 2.9: The running couplings of the three forces in the Standard Model (left) and the Minimal Supersymmetric Standard Model (right), taken from [39].

2.4 Proton decay

2.4.1 Theory

In the Standard Model, baryon and lepton number are conserved. The proton, being the lightest baryon, is therefore stable. GUTs introduce operators which violate baryon and lepton number, conserving $B - L$. This gives rise to proton decay. The mechanism for proton decay depends on the details of the model, as does the dominant final state. In non-SUSY GUTs, the fastest allowed proton decay is mediated by X bosons and suppressed by M_{GUT}^{-2} . The decay mode is $p \rightarrow e^+ \pi^0$ as shown in Figure 2.10a. The experimental bound at 90% confidence from Super-Kamiokande of 1.7×10^{34} years [21] requires a GUT scale higher than what the simplest models predict.

With the addition of SUSY, the $p \rightarrow e^+ \pi^0$ mediated by an X boson is still




Figure 2.10: (a) $p \rightarrow e^+ \pi^0$ mediated by an X boson. (b) $p \rightarrow K^+ \nu$ through a loop of supersymmetric particles.

allowed, but the predicted lifetime is much higher because the GUT scale increases by about a factor of 10. Since $\tau \propto M_X^4$, this increases the predicted lifetimes to order of $10^{34} - 10^{35}$ years, consistent with the experimental bound. However, an additional decay mechanism is introduced, mediated by sparticles. Instead of being proportional to the GUT scale, this diagram has a factor of M_{SUSY}^2 , the mass of the supersymmetric particles. The SUSY-mediated decay requires a second- or third-family quark in the final state [40, 41]. The dominant mode is predicted to be $p \rightarrow K^+ \nu$. Theoretically the neutrino is in a muon or tau flavor eigenstate. The flavor subscript is omitted here for consistency because experimental bounds obviously do not depend on the neutrino flavor.

Numerous different versions of GUTs exist, all of which predict proton decay but with differences in the dominant channel and lifetime [42]. The original GUT model is due to Georgi and Glashow [6], and predicts $p \rightarrow e^+ \pi^0$ with a lifetime of 10^{30-31} years. The simplest SUSY GUTs are due to Dimopoulos and Georgi [11], and to Sakai and Yanagida [12]. They predict $p \rightarrow K^+ \nu$ with $\tau \sim 10^{28-32}$ years [43].




Figure 2.11: Experimental limits for proton decay in different channels, along with ranges of theoretical predictions. Taken from [44].

Extensions to these simple models give longer lifetime predictions. Numerous SUSY GUTs predict $p \rightarrow K^+\nu$ with lifetimes of 10^{32-35} years [13, 14, 15, 16, 17, 18]. Other GUT models with and without SUSY predict $p \rightarrow e^+\pi^0$ with longer lifetimes ranging from 10^{34-39} years [7, 8, 9, 10]. These extended models are still consistent with experimental bounds.

2.4.2 Experimental searches

Many experiments have searched for proton decay in many different channels, though the most common searches are $p \rightarrow e^+\pi^0$ and $p \rightarrow K^+\nu$ due to theoretical motivations. No proton decay has been observed (Super-K has two candidates for $p \rightarrow \mu^+\pi^0$ with a background expectation of 0.8 [21]). A summary of experimental limits is shown in Figure 2.11. The Figure does not include Super-Kamiokande's Run IV data. The only significant difference is in $p \rightarrow K^+\nu$, where the limit has increased to 5.9×10^{33} years [22].

Super-Kamiokande has set the world's best limit in nearly every proton decay channel. Proton decay searches in Super-K look for events originating in a 22.5 kiloton volume of water in the center of the 50 kiloton tank. Cherenkov light from particles traveling faster than the speed of light in water ($0.75c$) is collected by photomultiplier tubes lining the interior. Particles emit a cone of light, which is reconstructed as a ring on the wall of Super-K. Muons lose energy relatively slowly, and travel in a straight trajectory in the absence of a magnetic field. This results in a filled-in ring that is smooth around the edge. Electrons lose energy more quickly by radiative processes, resulting in a hollow, fuzzy ring. Examples of e -like and μ -like rings are shown in Figure 2.12.

The Super-Kamiokande analysis technique for $e^+\pi^0$ along with several other meson-antilepton final states is described in detail in Refs. [46, 47]. The selection requires two or three e -like Cherenkov rings, with no other rings in the event. For three-ring events, the π^0 mass is reconstructed and required to be between 85 and 185 MeV/c^2 . Events with delayed “Michel electrons” from $\mu \rightarrow e$ or $\pi \rightarrow \mu \rightarrow e$ are excluded. The invariant mass of the proton is reconstructed from the rings and required to be between 800 and 1050 MeV/c^2 . With this selection, 45% of $p \rightarrow e^+\pi^0$ events are accepted, but none are observed in data.

The Super-K limit for $p \rightarrow K^+\nu$ is nearly a factor of three lower for the proton lifetime [22]. This is because of the difficulty of detecting kaons in a water Cherenkov detector. The K^+ from proton decay is well below the Cherenkov threshold, and never emits Cherenkov light. Events can be reconstructed only by observing the decay products of the K^+ , and the residual nucleus.

Super-K uses three different techniques to search for $p \rightarrow K^+\nu$. The first technique requires a triple coincidence in time of the de-excitation photon, μ^+ from




Figure 2.12: Rings of Cherenkov light recorded in Super-K due to a muon (top) and an electron (bottom). Taken from [45].

$K^+ \rightarrow \mu^+$ decay, and electron from $\mu^+ \rightarrow e^+$. This selection has the lowest predicted background rate of 1.11 events per Megaton-year. However, the signal efficiency is just 8.4%, primarily because of the low efficiency (19%) to reconstruct the prompt gamma ray from the de-excitation of ^{15}N . In its 260 kiloton-year exposure, Super-K observed 0 events with a background expectation of 0.24 [22]. The largest single contribution to the background is from NC K^+ production by atmospheric neutrinos at 48% of the total background, with 25% coming from CC ν_μ interactions.

The second analysis relaxes the requirement of a de-excitation photon. The efficiency improves to 43%, but the background expectation increases. Without the de-excitation photon, any CC ν_μ interaction where the muon happens to have approximately 236 MeV/ c of momentum will be signal-like. A fit is performed to the region between 215 and 260 MeV/ c and the data is consistent with simulated background, as shown in Figure 2.13.

The third analysis looks for $K^+ \rightarrow \pi^+\pi^0$. One or two e -like rings are required, with no other rings in the event. The reconstructed invariant mass of the π^0 must be between 85 and 185 MeV/ c^2 , with the π^0 momentum between 175 and 250 MeV/ c . Finally, there must be energy backward from the π^0 consistent with a low-momentum π^+ that is above Cherenkov threshold but below the threshold to reconstruct a ring. The simulated distribution of Cherenkov light as a function of the angle to the π^0 is shown in Figure 2.14. The upper plot shows the full 180°, where the larger peak is due to the π^0 and the small peak centered at 23° is the Cherenkov light from the π^+ , which is emitted at a 23° for π^+ at the momentum from K^+ decay. The lower plot zooms in on the light from the π^+ , and is overlaid on top of background simulation, which is mostly π^0 production by neutrinos.

A Michel electron from $\pi^+ \rightarrow \mu^+ \rightarrow e^+$ is also required. This selection is 8.7%




Figure 2.13: The muon momentum spectrum for single muon events, primarily from charged-current atmospheric ν_μ and $\bar{\nu}_\mu$ interactions. The red boxes show the allowed range of the atmospheric neutrino simulation. The black points represent data with statistical uncertainty. The blue histogram is what would be expected from an additional component due to $p \rightarrow K^+ \nu$, as muons from K^+ decay would give a spike at 236 MeV/c. The data is consistent with the atmospheric neutrino simulation without proton decay. Taken from [22].




Figure 2.14: Top: the simulated signal distribution for $K^+ \rightarrow \pi^+ \pi^0$, where the larger peak is due to the photons from π^0 decay, and the bump at 23° is the light from the π^+ . Bottom: Zoom in on the π^+ region, overlaid on the simulated atmospheric neutrino background simulation. Taken from [22].

efficient, with 2.09 background events expected per Megaton-year. No events were observed in data for the 260 kiloton-year exposure, with a background expectation of 0.44. Background for this selection is dominated by CC π^0 production with a low momentum muon (38%). K^+ production is the next largest contributor (37%).

Super-K has PMTs lining approximately 40% of its inner wall, while the coverage

Particle	Mass (MeV/ c^2)	Momentum (Mev/ c)	Kinetic energy (MeV)
e^\pm	0.5	0.6	0.3
μ^\pm	105.7	120	54
π^\pm	139.6	158	71
K^\pm	493.7	558	252
p	938.3	1060	477

Table 2.2: Cherenkov threshold in water for different particle species.

will be 20% for the much larger Hyper-K. In 2001, an accident damaged PMTs at Super-K, and the subsequent data taking period, called Run II, went forward with 20% photocoverage [22]. Proton decay signal efficiency dropped about 25% for the two low-background analyses, while background rates increased by factors of 2.5 and 1.5 for the $\mu^+\nu_\mu$ and $\pi^+\pi^0$ selections, respectively. Hyper-Kamiokande [23] will likely have photocoverage closer to that of Run II. Analysis techniques will be similar to those developed for Super-K, with background rates on the order of a few events per Megaton-year.

The neutrino K^+ production backgrounds for water Cherenkov detectors are events with a K^+ below Cherenkov threshold, no other final-state particles above threshold, and no particles that produce Michel electrons besides the one from $K^+ \rightarrow \mu^+ \rightarrow e^+$. The minimum momentum for particles to produce Cherenkov light in pure water is given in Table 2.2. Particles below the momentum given will be totally undetectable unless they decay into a particle above threshold. The efficiency for detecting a particle rises as the momentum increases. The threshold for ring detection is approximately 50 MeV/ c greater than the Cherenkov threshold. Not listed in the table is the π^0 , which decays promptly to $\gamma\gamma$ and is always detected. Few MeV photons have low efficiency, for example the efficiency for detecting the 6 MeV photon from ^{15}N decay is 19% [22].

Michel electrons are detected at high efficiency, enabling Super-K to reject events with π^+ or μ^+ in the final state. The capture rate for negatively charged particles is higher, and $\approx 100\%$ for π^- and K^- . The nuclear capture releases nucleons but not photons, so these events are not observed. The most common NC K^+ channel is $\nu p \rightarrow \nu K^+ \Lambda$, which is invisible when the Λ decays to $p\pi^-$ and both are below threshold. The Λ decay to $n\pi^0$ will cause the event to be rejected because of the π^0 decay. Another process that can be invisible for Super-K is $\nu n \rightarrow \nu K^+ \Sigma^-$ when the π^- from Σ^- decay is below threshold.

Muon-flavor charged-current events at low neutrino energy can also fake the $p \rightarrow K^+ \nu$ signal. In particular, single kaon production, $\nu_\mu N \rightarrow \mu^- K^+ N$, is invisible when the μ^- is below threshold and captures. Charged-current $K^+ \Lambda$ production is less likely to be invisible, as both the μ^- and the decay products of the Λ must be below Cherenkov threshold, and the minimum neutrino energy for such a process is higher and more likely to produce higher-energy muons.

The DUNE far detector will also search for proton decay [26]. At 40 kilotons of total mass, it cannot compete with Hyper-K in exposure, as a 10-year run would put the total exposure on par with Super-K and far below what might be expected from Hyper-K. However, DUNE is able to measure the K^+ momentum, and with essentially no threshold for charged hadrons it is able to detect low-energy recoiling nucleons that could indicate that the K^+ is due to an atmospheric neutrino interaction rather than a proton decay. If Hyper-K has an efficiency of $\sim 6\%$ like Super-K Run II, DUNE can make up for its factor of 10 deficit in mass with a factor of 10 increase in efficiency, perhaps without allowing more background.

Neutrino-induced backgrounds in DUNE could come from a process like $\nu n \rightarrow \nu K^+ \Sigma^-$ followed by $\Sigma^- \rightarrow \pi^- n$ where the π^- is very low energy and the neutron does

not interact in the detector. Another source is K_L^0 produced outside the detector which enter and then charge exchange by $K^0 p \rightarrow K^+ n$. These K_L^0 could come from atmospheric neutrino interactions in the rock above the detector, or from muon-nucleus interactions involving high-energy muons, for example $\mu n \rightarrow \mu K_L^0 \Lambda$. The 1,480 meters of rock overburden shields the detector from cosmic muons but doubles as a large target for atmospheric neutrino interactions that produce more muons.

2.5 Neutrino interactions

The simplest neutrino interaction is scattering off an atomic electron, $\nu e \rightarrow \nu e$ as in Figure 2.15. This process is point particle scattering and can be calculated precisely in electroweak theory. This interaction has a very important history in particle physics. In 1973, it was muon antineutrino-electron scattering which established the existence of the weak neutral current [31]. For electron neutrinos only, there is an additional contribution from charged-current scattering, where a W instead of a Z is exchanged as in Figure 2.15b. This was used by the Sudbury Neutrino Observatory to demonstrate that electron neutrinos from the sun were in fact transforming to muon and tau neutrinos [34, 35]. The fact that the cross section is known theoretically makes neutrino-electron scattering useful as a flux constraint in experiments, which is discussed in Section 3.2.4.

The cross section for neutrino-electron scattering is about 2,000 times smaller than that of neutrino-nucleon scattering. Nucleons are not point particles, complicating the interaction. Nuclei are even more complicated. Experiments studying neutrino oscillations need large detectors made of heavy, cheap nuclei, such as the carbon of NOvA [48], the oxygen of Super-Kamiokande [20], or the iron of MI-




Figure 2.15: (a) Neutral-current neutrino-electron scattering, equal for any neutrino flavor (b) Charged-current neutrino-electron scattering, for electron neutrinos only.

NOS [49]. The complicated strong interactions in the nuclear environment are at energies too low calculate with perturbative QCD, and effective theories must be used.

Neutrino-nucleon scattering is typically broken up into three categories: elastic (or quasi-elastic) scattering, resonance production, and deep inelastic scattering. When the squared four-momentum transfer to the nucleon, Q^2 , is small, the interaction can be approximated as scattering off the nucleon. Neutral-current elastic scattering (NCE), $\nu N \rightarrow \nu N$, is sometimes called quasi-elastic when the initial nucleon, N , is bound inside a nucleus. In, charged-current quasi-elastic scattering (CCQE), $\nu_\mu n \rightarrow \mu^- p$, the reaction is not truly elastic because the final-state lepton is charged, and because the target nucleon must be removed from the nucleus. CCQE is shown in Figure 2.4. CCQE is described theoretically by Llewellyn Smith in terms of form factors, which describe the structure of the nucleon and are taken from data [50].

At higher momentum transfer, the nucleon can be excited into a baryonic reso-




Figure 2.16: (a) Charged pion production through an intermediate Δ^{++} resonance. (b) Neutral pion production through an intermediate Δ^+ resonance.

nance. The most common excitation is to the $\Delta(1232)$ state, where the number in parentheses is the mass of the resonance in MeV/c^2 . The $\Delta(1232)$ is composed of up and down quarks, and can have any integer charge between $-1e$ and $2e$. It differs from a regular proton or neutron because the spin of the three quarks is aligned such that the total spin is $3/2$ instead of $1/2$. The Δ typically decays into a nucleon and a pion: $\Delta^{++} \rightarrow p\pi^+$, $\Delta^+ \rightarrow n\pi^+$ or $\Delta^+ \rightarrow p\pi^0$, $\Delta^0 \rightarrow n\pi^0$ or $\Delta^0 \rightarrow p\pi^-$, and $\Delta^- \rightarrow n\pi^-$. These events are called “resonant pion production” because the invariant mass of the pion-nucleon system will be peaked at the resonance mass. Heavier resonances are also produced, such as $\Delta(1600)$ and $\Delta(1700)$ [5]. The resonance widths become larger at higher masses, smearing out the resonance structure of the invariant mass spectrum. Diagrams for resonance production are shown in Figure 2.16. Resonances can also be produced via the neutral current in a similar way.

Heavier resonances can also decay into a nucleon and multiple pions, or into a kaon and strange baryon. There are also baryon resonances with strange quarks,

called Λ and Σ resonances. The lightest strange resonance, $\Lambda(1405)$, decays into a pion and Σ baryon. Heavier strange resonances decay into a nucleon and antikaon, either K^- or \bar{K}^0 . The strange resonance contribution to K^+ production is negligibly small because the K^+ has an antistrange quark ($S = 1$) instead of a strange quark ($S = -1$). There would be a large strange resonance contribution to K^+ production in a detector made of antimatter, however.

At high Q^2 , the probe wavelength is short, and the neutrino penetrates the nucleon and scatters off an individual quark. The struck quark undergoes hadronization, as described in Section 2.1.2, where pions and kaons are created. Kaons arise when gluons pair produce $s\bar{s}$. The \bar{s} quark combines with a d or u quark to form a K^0 or K^+ , respectively. The s quark can join the two spectator quarks in the struck nucleon to form a Λ or Σ baryon, or another antiquark to form a K^- or \bar{K}^0 meson. Production of pairs of strange particles, either by resonance decay or DIS, is called “associated production.” The total strangeness is zero in associated production, which is also called $\Delta S = 0$, meaning that the number of strange quarks is equal to the number of strange antiquarks. Associated K^+ production pairs can be kaon-baryon ($K^+\Lambda$, $K^+\Sigma^-$, or $K^+\Sigma^+$), or kaon-kaon (K^+K^- or $K^+\bar{K}^0$). Kaon production in the GENIE simulator is described in Sections 5.1.1 and 5.1.2. Examples of K^+ production diagrams by resonance decay and DIS are shown in Figure 2.17.

Neutrino interactions are best understood at low energies ($E_\nu \lesssim \text{few } 100\text{s MeV}$) where the scattering is predominantly elastic, and at high energies ($E_\nu \gtrsim 10 \text{ GeV}$) where the scattering is predominantly deep inelastic. Unfortunately, the so-called “transition region” from 1-10 GeV, where elastic, resonant, and deeply inelastic processes all contribute and neither low-energy nor high-energy approximations are valid, is of great importance. Neutrino oscillation searches looking for appearance




Figure 2.17: (a) Charged kaon production through an intermediate Δ^+ resonance. (b) Charged kaon production in deep inelastic scattering.

must be well above the 110 MeV threshold for charged-current ν_μ interactions. Because the larger Δm^2 is of order 10^{-3} eV^2 and the oscillation length depends on L/E , going to 10s of GeV in energy requires experimental baselines that are of order 10^4 km, greater than the radius of the earth. The K^+ production threshold is 0.8 GeV, so the relevant energy range for proton decay backgrounds is also in the transition region.

So far, this section has discussed neutrino-nucleon or neutrino-quark scattering, and has neglected to mention the nucleus. The nucleus is notorious for its complicating affect on neutrino scattering [51]. The first set of complications are initial-state effects. The nucleons bound inside the nucleus are not at rest, but are moving due to Fermi motion. A simple treatment of the nucleus still in use today is the Relativistic Fermi Gas model of Smith and Moniz [52], which treats the nucleons as quasi-independent, interacting only with the mean field of the nucleus. The nucleon momentum spectrum is given by a Fermi distribution. It neglects effects due to correlations between nucleons.

One effect of nucleon-nucleon correlations is a modification to the Fermi momentum. When two nucleons are interacting, they can have large equal and opposite momentum such that the net momentum of the pair is below the Fermi cutoff. This modification to the nucleon initial momentum is modeled by Bodek and Ritchie [53, 54]. When a neutrino scatters off a nucleon in a correlated pair, its partner can be ejected. This has been observed in electron-nucleus scattering [55], where 20% of nucleons in carbon are found to be in a correlated pair, 90% of which are np . An example from electron scattering is shown in Figure 2.18. One class of correlation is called “meson exchange currents,” where nucleons exchange a virtual pion inside the nucleus. One popular model of the lepton side of this process for neutrino energies below 10 GeV and events with zero pions is Gran *et al.* [56], and there are now predictions for the hadronic side as well [57]. This type of nuclear modeling is a very active area of theoretical work, and incorporating these improvements into neutrino simulations is important.

After the neutrino interaction occurs, any hadronic products must exit the nucleus, and may undergo FSI. The strong forces in the many-body system of the nucleus cannot be calculated exactly. Instead, the typical treatment is an intranuclear cascade, where hadrons are gradually stepped through the nucleus. At each step, possible interactions are considered, with interaction probabilities often coming from parameterizations of hadron-nucleus scattering data. These interactions include elastic scattering, knockout of a nucleon, meson absorption by the nucleus, and charge exchange.

Kaon-nucleus and pion-nucleus reactions differ because of strangeness conservation. Absorption is the dominant feature in the pion-nucleus inelastic cross section at pion kinetic energies in the few 100s of MeV. In K^- -nucleus scattering, the K^- can




Figure 2.18: An electron scatters off a proton in a correlated pair, ejecting the proton and its partner. Taken from [55].

be absorbed, converting a bound nucleon into a hyperon (Λ or Σ). The analogous process for K^+ -nucleus scattering is forbidden because there are no antibaryons in the nucleus. A K^+ produced inside the nucleus will exit unless it charge exchanges to a K^0 . In addition, K^+ can be produced in π^+ -nucleus reactions by strong processes such as $\pi^+ n \rightarrow K^+ \Lambda$. In the Giessen Boltzmann-Uehling-Uhlenbeck (GiBUU) model [58], this kind of reaction gives an enhancement to the K^+ production cross section at low K^+ momentum. In GENIE [59], the event generator used by MIN-

ERvA and many other experiments, 13% of K^+ produced in carbon reinteract before exiting the nucleus, distorting the spectrum toward lower kaon energies. GENIE does not include K^+ production either by pions or charge exchange in its FSI model.

2.6 Previous measurements of kaon production

Strange particle production by neutrinos has previously been observed in bubble chambers [60, 61, 62, 63, 64, 65, 66, 67, 68, 69, 70] and the straw tube detector NOMAD [71]. Bubble chambers have exquisite position resolution, and are fantastic detectors for Λ and K_S^0 , neutral particles which travel a few centimeters before decaying. When they decay to charged particles, $\Lambda \rightarrow p\pi^-$ or $K_S^0 \rightarrow \pi^+\pi^-$, the V-shaped decay is isolated from the neutrino interaction point. In a magnetic field, the oppositely-charged hadrons will bend apart as they propagate. An event like $\nu n \rightarrow \nu K_S^0 \Lambda$ produces a double-V pattern in the detector as shown in Figure 2.19, a trace from the ANL 12-foot bubble chamber [67].

Charged kaons are more difficult to detect in bubble chambers. If the K^+ is contained, it can be identified based on its decay. Otherwise, its ionization profile can be measured, but often cannot be distinguished from a pion or proton track. In events with a positively-identified Λ and an additional hadron track, the K^+ can be inferred. The ANL bubble chamber measured cross section ratios on deuterium

$$\begin{aligned}
 \frac{\sigma(\nu N \rightarrow \mu^- K^+ \Lambda)}{\sigma(\nu N \rightarrow \mu^- N \text{pions})} &= 0.04 \pm 0.03 \\
 \frac{\sigma(\nu N \rightarrow \mu^- K^+ N)}{\sigma(\nu N \rightarrow \mu^- N \text{pions})} &= 0.03 \pm 0.03 \\
 \frac{\sigma(\nu N \rightarrow \nu K_S^0 \Lambda)}{\sigma(\nu N \rightarrow \mu^- K^+ \Lambda)} &= 1.5 \pm 1.5
 \end{aligned} \tag{2.13}$$




Figure 2.19: The neutrino interaction $\nu n \rightarrow \nu K_S^0 \Lambda$ in the ANL 12-foot bubble chamber produces a double-V in the detector [67]

based on 7 strange particle events [66]. The BNL 7-foot bubble chamber also measured deuterium ratios

$$\begin{aligned} \frac{\sigma(\nu d \rightarrow \mu^- K_S^0 X)}{\sigma(\nu d \rightarrow \mu^- X)} &= (2.4 \pm 0.9)\% \\ \frac{\sigma(\nu d \rightarrow \mu^- \Lambda X)}{\sigma(\nu d \rightarrow \mu^- X)} &= (2.8 \pm 1.0)\% \end{aligned} \tag{2.14}$$

as well as the CC/NC ratio

$$\frac{\sigma(\nu d \rightarrow \nu V^0 X)}{\sigma(\nu d \rightarrow \mu^- V^0 X)} = 0.22 \pm 0.14 \quad (2.15)$$

where V^0 is any neutral strange particle, Λ or K^0 . The BNL measurement is based on 36 total strange events [68]. These experiments made the first observations of strangeness production by neutrinos and established the existence of different channels, but lack the statistics to be used as constraints to proton decay experiments.

2.7 Constraining proton decay backgrounds with MINERvA

The next generation of proton decay searches will have discovery potential for $p \rightarrow K^+ \nu$ at lifetimes above 10^{34} years, in the range predicted by SUSY GUT models. With the tenfold increase in exposure, Hyper-Kamiokande will have a background expectation of a few events. The background due to K^+ production by atmospheric neutrinos comes from theoretical models for neutrino cross sections that have not been directly tested by experiment. The MINERvA experiment combines a sensitive detector (Chapter 4) with the most intense muon neutrino beam ever produced (Chapter 3). For Hyper-K, the neutral current background is most important, especially for events with a K^+ and very little other energy in the final state.

The ideal measurement would be the cross section on a water target for K^+ production with no final-state particle above Cherenkov threshold, with the flux reweighted to the atmospheric neutrino spectrum. This is not possible for several

reasons. First, while MINERvA does have a water target, its mass is just 452 kg or 9% of the mass of the tracker fiducial volume, and reconstructing a complicated final state in the nuclear targets region of the detector is difficult. Instead, measurements on a hydrocarbon target (CH) are presented. Second, at low energies it is difficult to distinguish between photons from π^0 decay and protons ejected from nuclei by scattering neutrons. Third, the ability to reconstruct the energy of pions and protons depends on interactions inside the detector, which are different for water and CH . It is difficult to distinguish protons from pions at energies near the pion Cherenkov threshold. The active tracker of MINERvA is simply not well suited to measure Cherenkov thresholds in water. Fourth, the neutrino energy cannot be reconstructed in NC events, and the integrated flux of MINERvA is not the same as that of atmospheric neutrinos. While MINERvA has a constrained prediction of the flux spectrum, the neutrino energy is not known on an event-by-event basis.

MINERvA can still contribute to proton decay background constraints for Super-K and Hyper-K. Currently, cross section models are being used to make the prediction. A measurement of the flux-integrated differential cross section $d\sigma/dT_K$ for NC K^+ production on CH can be used to benchmark a prediction from simulation. A model can be used to extrapolate from the MINERvA flux to the atmospheric neutrino flux, from CH to H_2O , and from a semi-inclusive final state to one where no other particles are above threshold. The latter extrapolation is improved by a measurement of the visible energy in the detector for these events.

The K^+ FSI model is an important input for $p \rightarrow K^+\nu$ searches at DUNE. CC K^+ production in MINERvA is sensitive to the same FSI effects as a $p \rightarrow K^+\nu$ event. While the FSI models currently available in the event generators used in experiments have an incomplete treatment of kaons, comparing these predictions to

the shape of the K^+ kinetic energy spectrum measured by MINERvA is a way to demonstrate which missing features are most important. Adding a high-statistics measurement on a nuclei to existing deuterium data is a way to separate FSI effects from nucleon-level cross sections.

The NuMI neutrino beamline is discussed in Chapter 3. Chapter 4 describes the detector and calibration procedures, while the simulation used by MINERvA is described in Chapter 5. The reconstruction techniques used to isolate samples of CC and NC K^+ production events are in Chapter 6. The results are presented in Chapter 7. Chapter 8 is a separate third measurement, charged-current coherent K^+ production, which has never before been observed but is not relevant for proton decay.

Chapter 3

Neutrino Beam and Flux

Neutrinos for the MINERvA experiment come from the NuMI (Neutrinos at the Main Injector) beamline at Fermilab [72], currently the world’s most powerful neutrino beam. NuMI was originally designed for the MINOS long-baseline oscillation search, which began taking data in May, 2005. A diagram of the Fermilab accelerator complex, including the Main Injector, is shown in Figure 3.1. In this chapter, the production of neutrinos is described in Section 3.1, followed by the resulting neutrino flux prediction in Section 3.2.

3.1 The NuMI beamline

3.1.1 Proton acceleration

The Fermilab proton beam starts with an ion source, which magnetically removes H^- ions from hydrogen gas. A 750 kilovolt electrostatic field is applied, accelerating the ions to an energy of 750 keV. For the data used in the results presented in section 7, a Cockcroft-Walton generator was used for this step. In November 2012, this

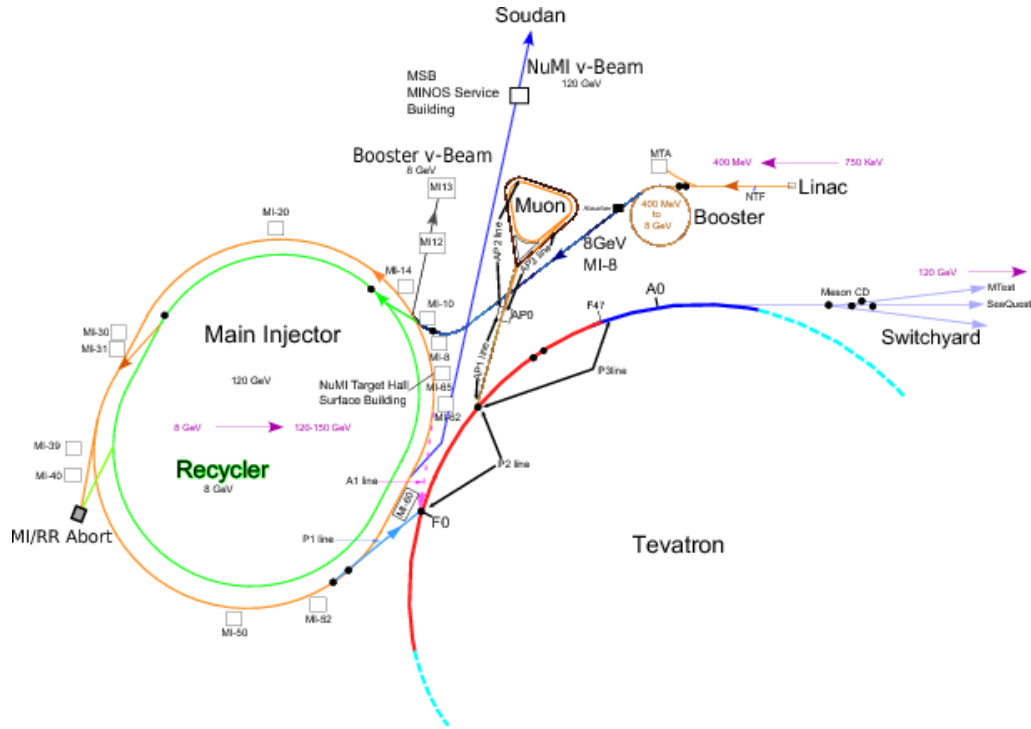


Figure 3.1: A diagram of the Fermilab accelerator complex, taken from [72].

room-sized accelerator was upgraded to a much smaller radio-frequency quadrupole (RFQ). The 750 keV ions are further accelerated to 400 MeV by the linear accelerator (Linac). They are then passed through a carbon foil, which removes the two electrons, resulting in a free proton beam. The protons are passed into the Booster synchrotron, which accelerates the beam to 8 GeV. The Booster acceleration is done in bunches with 53 MHz spacing, such that bunches are 18.5 ns apart. The bunches are grouped into $1.6 \mu\text{s}$ “batches.” Each batch corresponds to one revolution around the 474 meter accelerator ring.

The Booster batches are then fed into the Main Injector, a larger synchrotron with a circumference of 3319 meters. The Main Injector can accommodate six Booster batches, and accelerates protons to 120 GeV. Until 2011, one batch went to

the antiproton source when the Tevatron was operational. After the closing of the Tevatron, all six batches went to NuMI. The six batches lead to a beam “spill” that is $10\ \mu\text{s}$ in duration. A typical spill contains a total of between 2×10^{13} and 4×10^{13} protons.

3.1.2 The NuMI target

Protons from the Main Injector are bent downward 3.5° to account for the curvature of the Earth and point to the MINOS Far Detector facility in Soudan, Minnesota. The protons travel 350 meters to the NuMI target, a graphite rod 95 cm long. The target is designed to maximize the number of mesons produced by proton interactions in the target while minimizing the number of meson reinteractions. This is achieved by a long, thin target, with a 15×6.4 mm cross section. In the beam direction, the target is approximately two nuclear interaction lengths long such that 86% of protons will interact. Because of the small cross section, pions and kaons produced inside the target will typically exit out the side without themselves interacting inside the target. The target consists of 47 “fins” each 2 cm in length. A drawing of the NuMI target apparatus is shown in Figure 3.2.

3.1.3 Focusing system

Pions and kaons produced inside the target are focused by two magnetic horns, called Horn 1 (upstream) and Horn 2 (downstream). The focusing system is described in detail in Ref. [73]. Each horn has an inner and outer conductor, with 185 kA of pulsed current running in the direction approximately parallel to the proton beam in the inner conductor, and antiparallel to the beam in the outer conductor in




Figure 3.2: A drawing of the NuMI target and cooling system, taken from [72].

“forward horn current” (FHC) mode. The resulting magnetic field focuses positively-charged mesons in the proton beam direction, and bends negatively-charged mesons out of the beam. The π^+ decay in flight gives rise to a ν_μ -enriched beam, often referred to as “neutrino mode.” The direction of the horn current can be reversed, so that the horns focus negative mesons in so-called “reverse horn current” (RHC) or “antineutrino mode.”

Figure 3.3 shows a sketch of the focusing system, along with example particle trajectories. The solid blue π^+ is “overfocused” by Horn 1, and corrected by Horn




Figure 3.3: A drawing of the NuMI focusing system, showing example particle trajectories in Horn 1 and Horn 2. Taken from [72].

2. The dashed red π^+ trajectories represent “underfocused” and unfocused tracks in Horn 1. Mesons with very little transverse momentum off of the target will pass through the horn system totally unfocused, regardless of their charge. These mesons are typically very energetic, and give rise to high-energy neutrinos or antineutrinos.

The mesons are directed into a 675 meter decay pipe filled with low-pressure helium gas. Neutrinos (antineutrinos) are produced predominantly via the decay of π^+ (π^-) in FHC (RHC) beam mode: $\pi^+ \rightarrow \mu^+ \nu_\mu$ or $\pi^- \rightarrow \mu^- \bar{\nu}_\mu$. Electron neutrinos in the beam arise primarily from muon decays $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$ and $\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$, as well as from $K^+ \rightarrow e^+ \nu_e \pi^0$ and the charge conjugate.

So-called “wrong sign” particles are antineutrinos in FHC beam and neutrinos in RHC beam. The parent meson is of the wrong sign for the focusing horns. Often this is because the meson passes through the center of the two horns in the region

where there is no magnetic field. This occurs mostly for high-energy neutrinos (few 10s of GeV) because the parent mesons have very low transverse momentum and are typically the most energetic. They also arise from muon decays, for example in FHC mode a focused $\pi^+ \rightarrow \mu^+ \nu_\mu$ followed by $\mu^+ \rightarrow e^+ \nu_e \bar{\nu}_\mu$. In FHC mode, antineutrinos make up 11% of the total flux, but only 8% of the flux below 7 GeV. This is because the neutrinos in the “focusing peak” at 3.5 GeV come from pions that were bent by the horns. Very high energy pions that are produced at very low transverse momentum are not focused.

3.2 The neutrino flux prediction

The neutrino flux is predicted using a Geant4-based [74, 75] model of the target and focusing system, and is described in detail in Refs. [76, 77]. Proton interactions with the carbon target are simulated, and the resulting mesons are propagated through the horns and into the decay pipe. Each meson is weighted by the probability that the neutrino from its decay will intersect the MINERvA detector. The greatest challenge is predicting the meson spectra from the proton interactions in the target, and modeling subsequent interactions in the target as well as in other materials that mesons pass through prior to decay. Pion and kaon yields are predicted by Geant4 using the Fritiof-precompound [78, 79] and Bertini cascade [80, 81] models (FTFP_BERT) in bins of the Feynman scaling variable x_F and the meson transverse momentum p_T . Feynman x is defined as:

$$x_F \equiv \frac{p_{\parallel}}{p_{\parallel}^{max}} \simeq \frac{2p_{\parallel}}{\sqrt{s}}, \quad (3.1)$$

where p_{\parallel} is the meson momentum projected onto the proton beam axis, p_{\parallel}^{max} is the

maximum possible value of p_{\parallel} given the beam parameters, and \sqrt{s} is the center-of-mass energy. For inclusive high-energy scattering, cross sections expressed in terms of x_F are predicted to be largely independent of the incident beam energy [82].

3.2.1 The NA49 data

The Geant4 predictions are tuned to data from the NA49 experiment [83, 84]. NA49 measured pion yields in proton-Carbon interactions at 158 GeV for $p_T < 2\text{GeV}/c$ and $x_F < 0.5$. This data covers the relevant phase space for the neutrinos that pass through the MINERvA detector. The data is scaled to the NuMI beam energy of 120 GeV using the FLUKA simulation [85]. An uncertainty on the scaling is evaluated by scaling NA49 data at 158 GeV to 31 GeV and comparing the result to independent data from the NA61 experiment [86]. A weight, w_{NA49} , is defined as

$$w_{NA49} = \frac{\sigma_{NA49}(158\text{GeV})}{\sigma_{Geant4}(120\text{GeV})} \frac{\sigma_{FLUKA}(120\text{GeV})}{\sigma_{FLUKA}(158\text{GeV})}, \quad (3.2)$$

where σ_{NA49} is the measured differential cross section in the NA49 experiment in a bin of x_F and p_T , σ_{Geant4} is the Geant4 predicted cross section in the bin using FTFP_BERT, and σ_{FLUKA} is the cross section in FLUKA at a given proton energy, all as a function of x_F and p_T . More details can be found in Ref. [87]. The weight is applied to all simulated neutrinos for which the meson parent x_F and p_T fall in the bin. Figure 3.4 shows the weights for neutrinos with π^+ parents, the majority of the flux in FHC mode.




Figure 3.4: Weights applied to simulated neutrinos with π^+ parents based on NA49 proton-carbon scattering data in bins of x_F and p_T .

3.2.2 Extending to other interactions

The simple weight given in Equation 3.2 is for a neutrino from meson decay, where the meson was produced from a single interaction in the carbon target of a 120 GeV proton. Neutrinos also arise from mesons which are the products of multiple interactions, for example, $pC \rightarrow pX \rightarrow \pi^+ X$, $pC \rightarrow \pi^+ X \rightarrow \pi^+ X$ or $pC \rightarrow nX \rightarrow \pi^+ X$. Also, reinteractions can take place outside the target in a material other than carbon. Hadron production data can be extended to cover these more complicated neutrino origins with simple theoretical arguments.

Protons that interact twice inside the target are constrained by an NA49 measurement of proton production in proton interactions in a thin carbon target. This data is used to reweight $pC \rightarrow pX$ interactions. The subsequent proton interaction

at an energy lower than 120 GeV is covered by the same pion data. The FLUKA scaling in equation 3.2 is evaluated to the energy of the reinteracting proton. Pion reinteractions are not corrected by data. A 40% uncertainty is applied to neutrinos from pions which reinteract before decaying.

For interactions in materials other than carbon, a power law is assumed. The exponent is parameterized as the product of second-order polynomials in x_F and p_T . The free parameters are determined by a fit to data for proton interactions in various nuclei [88]. Interactions in materials other than carbon are reweighted using the power law with an additional uncertainty between 2.5% and 15% depending on x_F and p_T .

Meson production in neutron interactions is included by an isospin symmetry argument. Because the target is made of an isoscalar nucleus,

$$\sigma(pC \rightarrow \pi^\pm X) \approx \sigma(nC \rightarrow \pi^\mp X), \quad (3.3)$$

and opposite-sign pion yields from pC data can be used to constrain nC reactions. Equation 3.3 becomes exact if electromagnetic interactions are ignored and the proton and neutron masses are taken to be equal. Neutron-induced meson production in nuclei other than carbon is negligibly small. It is reweighted using data from proton-carbon interactions, and extrapolated to other nuclei using the Geant model.

3.2.3 Flux and uncertainties

The predicted neutrino flux is shown in Figure 3.5, along with the ratio of the corrected to uncorrected flux. The total flux in bins of neutrino energy is given in Table 3.1 for muon neutrinos and antineutrinos. In the charged-current result

presented in Chapter 7, antineutrino-induced events are subtracted as a background, and only the ν_μ flux is used. However, in the neutral-current result, there is no final-state observable that distinguishes ν -induced and $\bar{\nu}$ -induced scattering on an event-by-event basis, and both are considered signal. Neutral-current scattering is flavor symmetric, and electron-flavor neutrinos are also included. The ν_e flux prediction is not shown but is approximately 1% of the total.

The correction to the Geant4 flux prediction in a bin of neutrino energy is essentially the average weight in Figure 3.4 for all neutrinos of that energy. Statistical and systematic uncertainties from the NA49 data are propagated onto the flux prediction. There are 11 parameters that are varied within their uncertainties to form 100 “universes” for each parameter, and a flux prediction in each bin of neutrino energy is calculated in each universe. The spread of the universes in each bin gives the uncertainty due to that parameter.

A summary of the hadron production uncertainties on the flux prediction is shown in Figure 3.6. The largest uncertainty, labeled “nucleon-A,” covers proton and neutron interactions with nuclei other than carbon. These interactions are constrained using carbon data and a power-law extrapolation that brings in an additional uncertainty. The solid blue line, labeled “ $pC \rightarrow \pi X$,” is dominated by NA49 uncertainties in the simplest-case weight given by Equation 3.2. The solid tan “meson-inc.” line is primarily pion reinteractions in the target. The solid green “target att.” covers an attenuation correction applying thin-target constraints to a target three interaction lengths long, and arises because the number of protons remaining as a function of distance depends on the interaction cross section.

The importance of kaon production ($pC \rightarrow KX$, not to be confused with $\nu C \rightarrow KX$ which is the primary topic of this thesis) rises with increasing neutrino energy as




Figure 3.5: The predicted neutrino flux after applying corrections from NA49 (top), and the ratio of the corrected to uncorrected predictions showing the effect of the NA49 weights (bottom).




Figure 3.6: A summary of fractional uncertainties on the flux as a function of neutrino energy.

can be seen by the relative uncertainties in Figure 3.6. Kaon yields are constrained by MIPP data [89], and K/π ratio measurements [90]. Statistical uncertainties on kaons from MIPP are 3 – 4 times larger than NA49 uncertainties on pion yields in the most important x_F - p_T region, but the contribution from kaons is small below neutrino energies of 10 GeV.

3.2.4 Neutrino-electron scattering constraint

Neutrino scattering off atomic electrons, $\nu e \rightarrow \nu e$ shown in Figure 2.15, is point particle scattering and can be calculated in electroweak theory. Because the cross section is known, a measurement of the rate of neutrino-electron scattering is used to constrain the prediction of the neutrino flux. The analysis technique is described

E_ν (GeV)	$10^{-9} \nu_\mu/\text{cm}^2/\text{POT}$	$10^{-9} \bar{\nu}_\mu/\text{cm}^2/\text{POT}$
0 – 1	1.0331	0.3027
1 – 2	4.3611	0.3575
2 – 3	7.4333	0.3670
3 – 4	7.9013	0.3446
4 – 5	3.2984	0.2963
5 – 6	1.2193	0.2649
6 – 7	0.7644	0.2427
7 – 8	0.5671	0.2116
8 – 9	0.4398	0.1797
9 – 10	0.3834	0.1478
10 – 15	1.1135	0.4273
15 – 20	0.4664	0.1519
20 – 25	0.1959	0.0564
25 – 30	0.0918	0.0233
30 – 35	0.0668	0.0116
35 – 40	0.0553	0.0065
40 – 45	0.0474	0.0033
45 – 50	0.0325	0.0022
50 – 60	0.0267	0.0021
60 – 70	0.0057	0.0005
70 – 80	0.0023	0.0002
80 – 90	0.0006	0.0000
90 – 100	0.0001	0.0000
100 – 120	0.0000	0.0000

Table 3.1: The predicted ν_μ and $\bar{\nu}_\mu$ (wrong sign) flux per POT in FHC beam mode.

in Ref. [91].

The flux constraint is applied by producing an electron spectrum for neutrino-electron scattering in each flux universe. The universe spectra are compared to the spectrum observed in data, and a weight is computed. The effect of the weight is to favor universes where the number of neutrino-electron events, and the shape of the electron energy spectrum, is consistent with the data. The distributions of integrated flux for neutrino energies from 2 to 10 GeV before and after applying the constraint are shown in Figure 3.8.




Figure 3.7: The flux prediction in each universe is weighted based on its consistency with the measured spectrum of neutrino-electron events.

The resulting flux is very similar to the prediction from hadron production. With the updates of Refs. [76, 77], the uncertainty on the hadron production constraint is smaller than the statistical uncertainty on the neutrino-electron measurement, and the $\nu - e$ constraint reduces the overall flux uncertainty by only about 10% of itself. The $\nu - e$ constraint is not used for the analysis presented in this thesis, or in the flux given in Table 3.1.

The flux and flux uncertainties before and after the neutrino-electron scattering constraint is applied are shown in Figure 3.8.

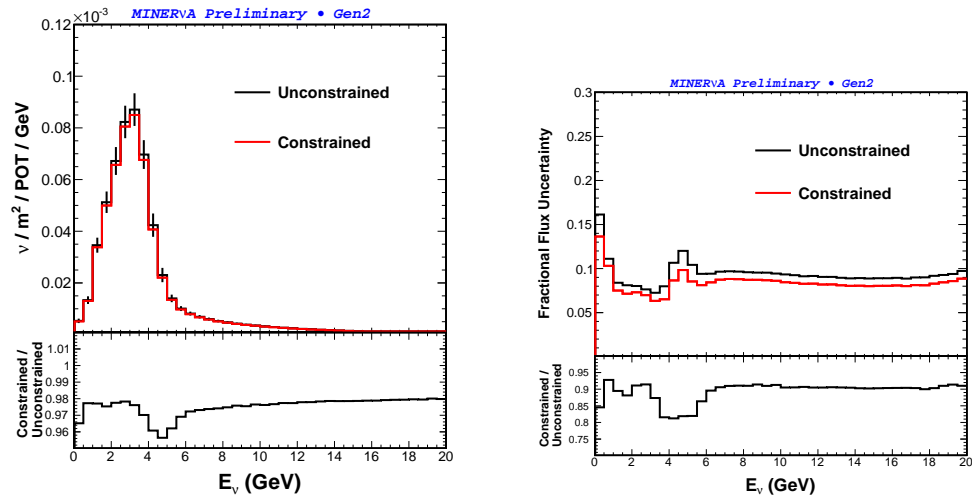


Figure 3.8: The flux prediction before and after the neutrino-electron constraint is applied (left), and the uncertainty on the flux prediction (right).

Chapter 4

The MINERvA Detector

The MINERvA detector [4, 92] consists of a core of solid plastic scintillator, surrounded by sampling calorimeters for containment. The central “tracker” region is fully-active. On the sides and downstream end are electromagnetic, and hadronic calorimeters, which alternate scintillator planes with passive planes of lead, and steel, respectively. Upstream of the tracker is a region containing targets made of carbon, water, iron, and lead, interspersed with scintillator. These targets, along with a cryogenic liquid helium target upstream of the detector, are used to compare neutrino interactions on different nuclei in the same beam. A schematic diagram of the detector is shown in Fig 4.1. The beam is angled downward at 3.5° to account for the curvature of the Earth in the MINOS long-baseline experiment. In the MINERvA coordinate system, the z axis is horizontal, 3.5° upward from the beam axis.

The plastic scintillator is made up of strips that are roughly two meters in length and are arranged into planes stacked perpendicular to the horizontal axis. The cross section of the detector is hexagonal, and the orientation of the strips is rotated 0°

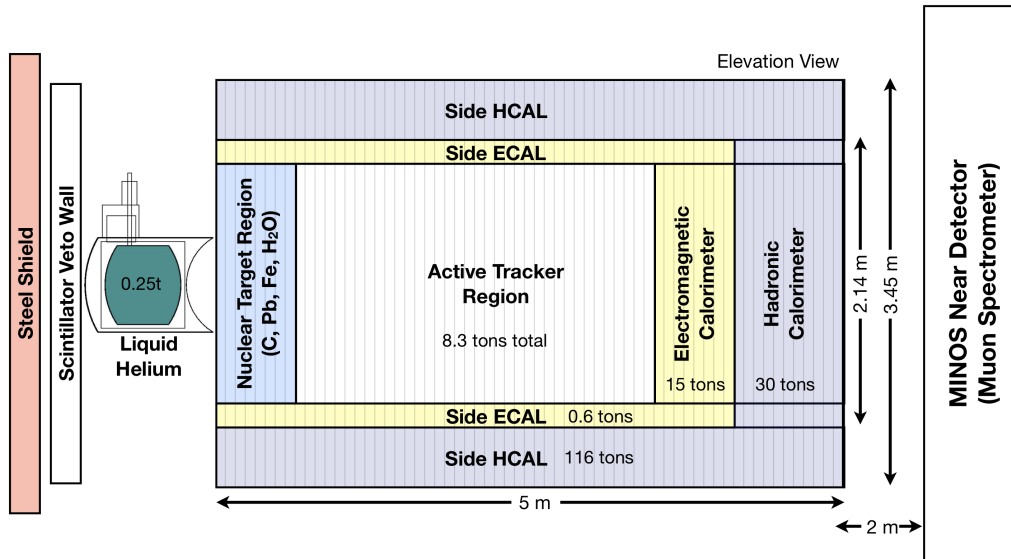


Figure 4.1: A diagram of the MINERvA detector viewed from above. The neutrinos are incident from the left.

and $\pm 60^\circ$ with respect to the vertical axis to enable unambiguous three-dimensional tracking of charged particles. The “X” orientation has vertical strips, with “U” and “V” orientations rotated $\pm 60^\circ$. The scintillator strip orientations can be seen in Figure 4.2. The planes are ordered UXVX throughout the detector so that there are twice as many X planes as U or V. The stacking of planes is shown in Figure 4.2.

Each plane contains 127 strips. The cross section of the strips is triangular with a base edge of 3.4 cm and a height of 1.7 cm. The triangles are overlaid such that a charged particle passing through a plane will deposit energy in two adjacent strips. The relative size of the energy deposits in the two strips is used to determine the position of the particle with a resolution of 3 mm. A photograph of the detector as




Figure 4.2: The stacking of scintillator planes (left), and a plane (right) with the three scintillator strip orientations. The pink collar is the side ECAL, and the cyan is the outer detector.




Figure 4.3: (Left) The detector as seen by the neutrinos. The inner hexagon is the “inner detector,” surrounded the side hadronic calorimeter. (Right) The cross section of a scintillator strip, and an example of how the strips are arranged into planes.

seen by the neutrinos, and the cross section of the scintillator strips are shown in Fig 4.3.

Upstream of each scintillator plane in the electromagnetic calorimeter is a 2 mm sheet of lead, corresponding to 0.35 electromagnetic radiation lengths. There are 20 such planes in the downstream calorimeter, as well as a lead collar that covers the outer 10 cm of each plane in the tracker. The 20 most downstream scintillator planes form the hadronic calorimeter, with a 1-inch steel plane preceding each plane of scintillator. The steel corresponds to 3 nuclear interaction lengths. The side hadronic calorimeter is located in the “outer detector.” In this region, four “stories” of scintillator strips run along the hexagonal edge, with steel in between the stories. The outer detector contains hadrons traveling transverse to the direction of the beam. It is also used to identify side-exiting muons, since hadrons typically do not traverse the outer detector without interacting. The side calorimeters can be seen in Figure 4.2

The MINOS near detector [49] is located 2 m downstream of MINERvA. MINOS is a magnetized iron calorimeter, and is used to measure the charge and momentum of muons exiting MINERvA. Muon momentum is measured to within 2% by either range (for muons up to 10 GeV) or curvature. The magnetic field also enables event-by-event muon charge identification, which is equivalent to discriminating between neutrino- and antineutrino-induced charged-current interactions.

4.1 Readout and Electronics

Scintillation light is absorbed and re-emitted by a wavelength-shifting optical fiber in the center of each strip. The wavelength shifter is a Kuraray S-35 multiclاد J-type with a 1.2 mm diameter, with Y11 fibers and 175 ppm dopant [92]. The fiber is mirrored at one end with an 80% reflection coefficient and measured by a Hama-

matsu R7600 64-channel multi-anode photomultiplier tube (PMT) [93] read out by a front-end board (FEB) at the other. Light incident on the photocathode produces photoelectrons (PE), which are converted to a charge by the PMT and digitized by one of six application-specific integrated circuit (ASIC) chips called TriP-t, which are mounted onto each FEB. A “hit” is a timestamped energy deposition in a single scintillator strip. When a charge called the “discriminator threshold” is exceeded, charge is integrated for 151 ns such that all subsequent energy depositions in the same strip will be grouped together. The uncalibrated hit time is the time of the earliest charge recorded in that strip, and is measured with a resolution of 2.2 ns due to the FEB.

Immediately following the integration window is a 151 ns period of “dead time” while the FEB is read out. During this period, no timing information is kept, and energy deposited does not generally result in charge being recorded. Dead time is important when two neutrino interactions occur close together in time and overlap spatially. Energy deposits resulting from the second neutrino interaction may not be observed if they occur on channels which are temporarily “dead” due to the readout of the first event.

The FEBs are linked together into chains of at most 10 boards. Up to four chains are read out by a chain read-out controller (CROC). The CROCs receive trigger and timing signals from the CROC interface module (CRIM). The data acquisition system and readout electronics are described in Ref. [94].

Calibrations are applied to convert ADC counts recorded by the FEB to energy. The timestamps of hits are also calibrated to correct for offsets between different FEBs, as well as PE-dependent “time slewing” due to the decay times of the scintillator and optical fiber. The conversion from ADC counts to fully-calibrated energy

is

$$E(s, t) = (\text{ADC}(s, t) - P(s, t)) \times F(s) \times \frac{1}{G(s, t)} \times A(s) \times S(s, t) \times M(t), \quad (4.1)$$

where s indicates that a calibration constant is determined separately for each strip, and t indicates that the correction is determined separately for periods of time ranging from days to weeks. E is the fully-calibrated energy used in physics analyses, ADC is the raw number of counts, P is the pedestal, F is the charge per ADC count for a given FEB, G is the PMT gain, A is an attenuation factor, S is a “strip-to-strip” constant that accounts for differences in response in different scintillator strips, and M is the “MEU” or muon energy unit factor, which sets the absolute energy scale based on muons.

The pedestal, FEB, gain, and attenuation calibrations are described in Section 4.2. The algorithm used to separate multiple neutrino interactions in the same beam spill is described in section 4.3. The strip-to-strip, MEU, and timing calibrations are described in detail in sections 4.4, 4.5, and 4.6, respectively. More details on the design and performance of the MINERvA detector can be found in Ref. [4].

4.2 Low-level calibrations

4.2.1 Pedestal calibration

The “pedestal” $P(s, t)$ is the dark current of the electronics, the level of ADC counts observed in the absence of energy in the detector. It is determined by reading out the detector for 16 μs in between beam spills when no activity is expected. An outlier




Figure 4.4: The ADC count distribution for a run with no beam activity is used to measure the pedestal. The outlier at 555 counts is rejected using an algorithm based on Peirce’s criterion.

removal scheme based on Peirce’s criterion [95] is used to reject signals well above the pedestal peak. These signals can be due to cosmic ray muons, electronics noise or radioactivity. Cosmic rays penetrate the 100 m overburden of MINERvA at a rate of about 18 Hz. A pedestal run of 750 readout gates occurs approximately once per day, meaning that a cosmic ray coincides with a $16 \mu\text{s}$ pedestal gate approximately once per week.

An example of a pedestal measurement with one outlier is shown in Figure 4.4. A typical pedestal is around 450 ADC counts. Measured pedestals vary by 7% across all channels, and are observed to be stable over time.

4.2.2 FEB calibration

The FEB constant $F(s)$ is measured for each board prior to installing the FEB on the detector. The bench measurement consists of injecting charges and recording the output ADC counts. The six TriP-t chips are divided into high, medium, and low gain channels. Four of the chips contain the high and medium gains for 16 channels apiece, while the remaining two chips have the low gains for 32 channels each. The medium gain is used when the high gain saturates, followed by the low gain when the medium channel saturates.

The FEB calibration is a trilinear fit to each of the three gain channels of each FEB. For each FEB, 18 parameters are stored: 3 slopes and the positions of the 2 kinks, for each of the 3 gain channels. An example of the digitized response for a typical FEB is shown in Figure 4.5.

4.2.3 PMT gain calibration

The PMT gain $G(s, t)$ is the charge per photoelectron. It is measured *in situ* by light injection, where light-emitting diodes are used to flash individual channels and the charge is recorded. The gain is the ratio of the measured charge to the number of PE times the electron charge, e . A typical gain is 6×10^5 , and the RMS over all channels in the detector is about 20%. Light injection data is taken once per day. A light injection gate occurs between beam gates, and about 5000 light injection gates are taken each day to perform the calibration. The statistical fluctuations for a single pixel from day to day are 3-5%.




Figure 4.5: The FEB calibration measures the ADC counts as a function of input charge for the three gain ranges.

4.2.4 Attenuation correction

The attenuation correction $A(s)$ is the product of three separate attenuation factors. The first is the “strip attenuation,” which corrects for attenuation in the fiber inside the scintillator strip. This factor is mapped out as a function of position along the strip using a radioactive source. The second factor corrects for attenuation in the wavelength-shifting fiber but outside the scintillator strip. This is called the “baggie attenuation” because the fibers are in a light-tight baggie as they come out of the scintillator plane. The wavelength-shifting fiber is coupled to a cheaper clear fiber, which is then optically coupled to the PMT face. The third factor corrects for attenuation in the clear fiber. The fiber routing can be seen in Figure 4.6.

The baggie and clear fiber attenuation are constant for a given strip because




Figure 4.6: A drawing of the optical fibers between the edge of the scintillator plane and the PMT, showing the “baggie” fiber (where the fibers are brought together), and the tighter clear fiber.

the light always travels the same distance in these sections of fiber. The baggie fiber length in particular varies significantly between strips, as strips on the edges of the plane typically have longer fiber lengths in the baggie. The strip attenuation correction depends on the location of the charged particle in the strip. This is known only when the hit is part of a three-dimensional reconstructed object. Otherwise, the location of the ionization energy along the strip cannot be determined, and the attenuation correction is applied to the center of the strip.

4.3 Time slicing

Due to the high intensity of the NuMI beam, a single spill typically produces multiple neutrino interactions that lead to energy in MINERvA. Most of these interactions

take place either in the rock upstream of the detector or in the outer detector. Timing information is used to separate hits into distinct “time slices,” where each slice ideally contains the products of a single neutrino interaction. The time slicing algorithm is run prior to any other reconstruction.

Hits are first sorted in time. An 80-ns window slides forward in time until the total pulse height of all hits in the time window exceeds 10 PE, at which point a time slice is opened. The window continues to slide forward, and hits are added to the slice until the total pulse height of the 80-ns window falls below 10 PE. A typical time slice is between 100 and 200 ns wide. The time of flight for a particle traveling at the speed of light through the detector is 10 ns, so the majority of the width is due to the spread in reconstructed times. The relationship between timing and pulse height is discussed in section 4.6.

An example of time slicing is shown in Figure 4.7. The top display shows an entire 10 μ s beam spill, while the bottom display shows a single time slice. In the histograms of hit times, the colors represent reconstructed time slices, and the hits shown in black are not part of any time slice because they occur at times such that no 80-ns window containing them has at least 10 PE. Such hits are excluded from reconstruction, and are typically noise due to pedestal fluctuations, or PMT afterpulsing from previous hits. The three panels of each display are different two-dimensional projections of the same event. They are surrounded by the six towers of the outer detector. The horizontal axis is the module number, and the vertical axis is the strip number within each scintillator plane.

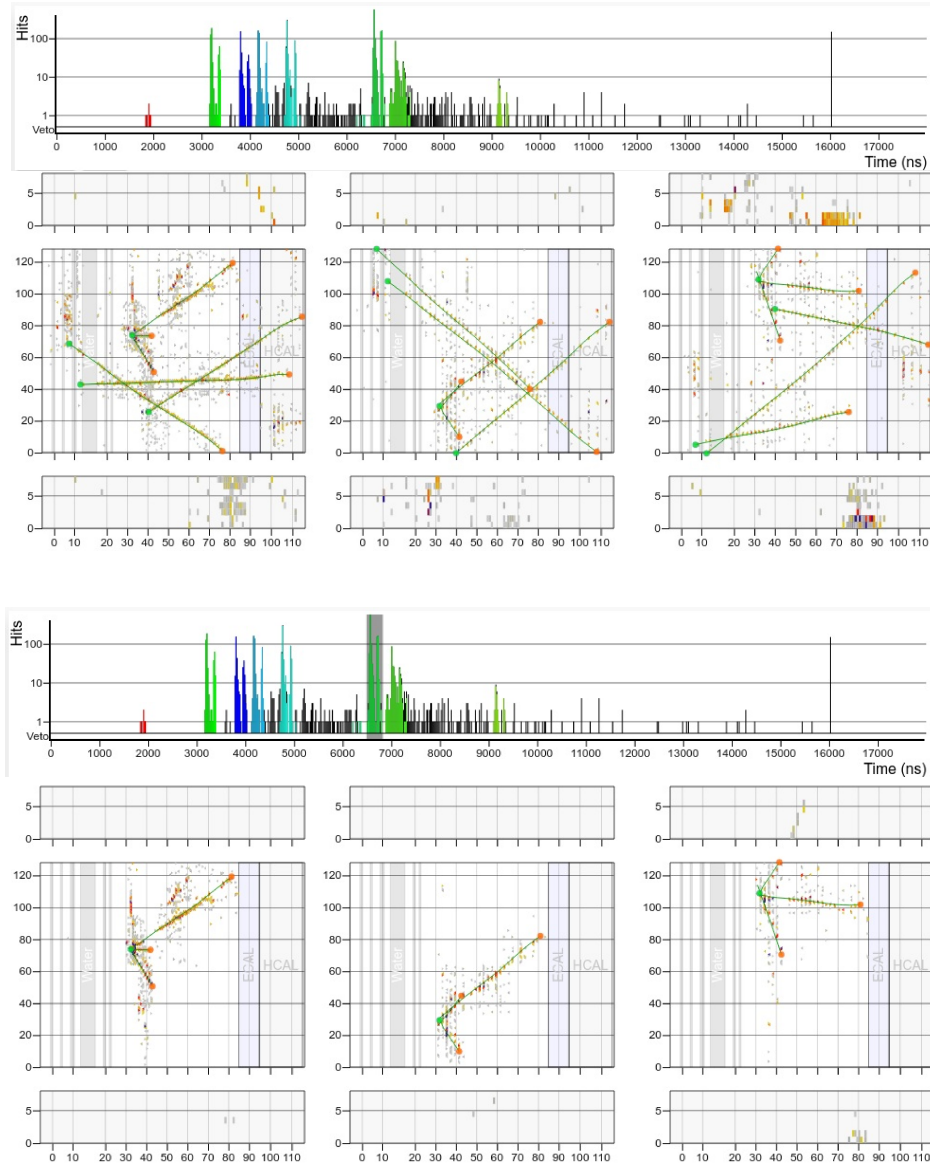


Figure 4.7: Top: An event display showing a full gate, with the histogram of hit times above. Bottom: The same event for a single time slice, with the hit times for the slice highlighted.

4.4 Strip-to-strip calibration

Variations in light level between inner detector strips are expected. These variations are caused by differences in the composition of batches of scintillator, air bubbles

in the epoxy used to fill the fiber hole, couplings between the optical fibers and photomultiplier tube, or some combination thereof. The goal of the strip-to-strip calibration is to correct the total effect of these components so that the energy response is uniform throughout the detector.

The strip-to-strip calibration consists of an alignment correction, a strip-level correction, a plane-level correction, and dead channel identification. In addition, it is used to identify errors in the mapping between electronics channels and physical detector strips. The procedure is based on reconstructing the energy deposited in a strip by a through-going muon, and the path length of the muon in that strip. The calibration is performed using a sample of “rock muons” that originate from neutrino interactions in the rock between the beam absorber and the detector hall. To be considered, a track must enter either the front or sides of the detector, and must also exit either the rear or sides. It must be the only reconstructed track in the time slice, and the amount of energy not associated to the track in the slice must not exceed 100 MeV to avoid accidental overlap due to pile-up. Muons are not required to be reconstructed in MINOS, so the momentum is not known. This allows the calibration to include muons which exit the side of MINERvA, which is important because it is necessary to have occupancy in every strip, including those at the edge of the detector. An example of a rock muon is shown in Figure 4.8.

For muons in the $0.1 - 10$ GeV range, the energy deposit per cm is largely independent of the muon energy. We require that the energy deposited by a muon be the same in every strip of the detector, and calculate a correction factor for each strip.

As strip-to-strip requires sufficient statistics to accurately determine the energy in each channel, a calibration table typically represents at least 300,000 muons,




Figure 4.8: A rock muon enters the front of the MINERvA detector and exits the rear. The green line is the reconstructed track. The green circle is the reconstructed start point, and the orange circle is the reconstructed end point. Each triangle is one hit, with the color proportional to the energy deposit.

about one month of data in normal running conditions during the NuMI low-energy era. The boundaries of these intervals of validity typically correspond to hardware changes in the detector. Due to the limitation of statistics, not every hardware change represents a new strip-to-strip interval of validity (IOV). Figure 4.9 shows the uncalibrated energy in each channel for one IOV.

The calibration corrects to the ratio of energy per unit path length. The path length calculation requires that the position of the planes is known to better than 1 millimeter in the direction transverse to the strips. Prior to calibrating the energy scale, a set of alignment corrections is determined for each plane.




Figure 4.9: The uncorrected truncated mean energy is plotted for each strip, expressed relative to the overall detector mean. The z-axis projection has a standard deviation of 19.1%.

4.4.1 Plane alignment correction

The alignment corrections are generated by comparing the energy in a strip to the point where the muon intersects the triangle base. In Figure 4.10, a hypothetical muon track is drawn intersecting two adjacent triangular strips. To account for the angle of the track, the energy is corrected for normal incidence by the multiplicative factor

$$C = \frac{\text{path}_{\text{normal}}}{\text{path}_{\text{actual}}} \quad (4.2)$$




Figure 4.10: A hypothetical rock muon intersects two adjacent strips.

where C is the correction factor, $\text{path}_{\text{normal}}$ is the normal-incidence path corresponding to the base position and $\text{path}_{\text{actual}}$ is the actual muon path length in the strip. In Figure 4.10, $\text{path}_{\text{normal}}$ ($\text{path}_{\text{actual}}$) is shown by a dashed (solid) line. Figure 4.10 shows a two-dimensional projection of the muon path. The actual path can contain a longitudinal component but the normal incidence path cannot.

The maximum normal-incidence corrected energy must occur where the normally incident path is longest. In a perfectly-aligned detector this will be the center of the strip. To determine the alignment shift, we look for the base position where the energy is maximized. The plane is treated as a rigid object, with all 127 strips shifted together. The procedure is sensitive only to shifts in the direction transverse

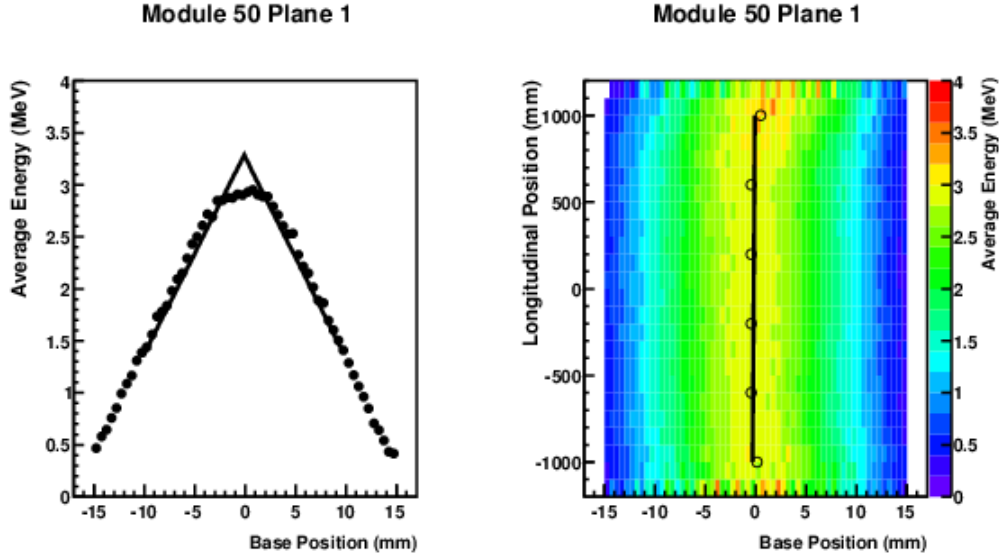


Figure 4.11: (Left) The average energy is plotted against position along the triangle base to calculate the total transverse shift. The rounded tip is due to the optical fiber in the center of the strip which does not scintillate. (Right) The shift is computed separately in slices of longitudinal position and fit to give the rotation.

to the strip direction. In addition to shifts, we correct for rotations about the z -axis of the detector. This is done by performing an identical fit in slices of longitudinal position. The resulting shifts are then fit linearly, and the rotation is given by

$$R = \arctan \frac{ds}{dL} \quad (4.3)$$

where R is the rotation in radians, L is the longitudinal position along a strip and s is the shift calculated at that position. The fits are shown in Figure 4.11.

The alignment correction is static and does not need to be recomputed for every interval of the strip-to-strip calibration. The resolution of the alignment is 0.3 mm and 0.5 mrad on average, and further iterations do not improve the resolution. After several years of data taking, the alignment was redone to test for shifts in the plane

positions over long time periods. It was found to be consistent with the detector not moving at all.

4.4.2 Plex and dead channels

After correcting the alignment, dead strips and errors in the “plex” map are identified. The plex maps electronics channels to physical detector strips, and problems occur when optical fibers are flipped during installation. Dead strips are those that report no energy, most commonly because the optical fiber is broken. In both cases, these issues are identified by looking for tracks that pass through a strip and leave no energy. For tracks with at least 2 mm of path length in a strip, about 5% will not deposit energy in a healthy strip due to Poisson fluctuations to zero photoelectrons. For dead strips, the percentage is much higher, often near 100%. If there is a fiber break along the strip, then it is possible to observe zero energy for some large fraction of tracks but not 100%.

Any strip where greater than 30% of tracks with at least 2 mm of path length result in no energy deposit is considered dead. This cut also removes a small number of strips where the energy response is extremely low, and Poisson fluctuations to zero PE are greater than 30%. These strips are not used in analysis in data or in simulation. Occasionally, the failure is in the electronics, and a new PMT can revive a strip. Approximately 60 of the 26,000 inner detector strips are dead. The distribution of the fraction of tracks with at least 2 mm of path length and zero energy is shown in Figure 4.12.

Plex map errors are identified by searching for pairs of strips where energy is observed in one strip when a track passes through another. When this happens consistently, it is an indication that the strips are swapped in the electronics, and




Figure 4.12: The fraction of tracks that result in zero energy deposit. Channels are considered dead if the fraction of tracks with zero energy is greater than 0.3. Healthy channels typically have zero energy 5-7% of the time due to Poisson fluctuations to zero photoelectrons, and the small peak near 1.0 is dominated by strips with broken or damaged optical fibers.

the map is updated. There are approximately 150 inner detector strips where the fiber is not coupled to the correct electronics channel. All of them are identified and fixed by the algorithm prior to running the strip-to-strip calibration.

4.4.3 Strip-level correction with the truncated mean

The strip-level correction forces the truncated mean energy per unit path length for through-going muons to be the same in every strip. The truncated mean is

used to avoid sensitivity to radiative processes and nucleon knock-out, which are dependent on the muon energy. The peak of the distribution of muon energy/cm is due to minimum ionization and does not depend on muon energy. Fitting for the peak requires a few thousand rock muons passing through a strip, while the truncated mean calculation is reliable with only a few hundred. This is especially important near the edges of the detector where there are fewer tracks. This allows the calibration interval to be shorter, making it easier to split intervals when FEBs are replaced, and improving the calibration of non-uniform scintillator degradation.

The truncated mean is computed iteratively. For the first iteration, the full mean of hits from 0 to 20 MeV per centimeter path length is used. In each subsequent iteration, the mean is taken considering events between 50% and 150% of the mean from the previous iteration. The procedure converges after three or four iterations but eight are used for redundancy. The calculated constants are normalized such that the average constant is exactly 1.0 over all healthy strips in the detector. The constant C_i for strip i is

$$C_i = \frac{\frac{1}{x_i}}{\frac{1}{N} \sum_j \frac{1}{x_j}} \quad (4.4)$$

where x_i is the truncated mean energy in strip i , N is the number of good channels in the inner detector and the sum in the denominator is over all good channels, indexed by j . Only tracks with at least 2 mm of path length in a given strip are used. An example of the energy per path distribution with the bounds of the truncated mean for five iterations is shown in Figure 4.13.




Figure 4.13: An example distribution of energy per path length. The truncated mean is computed iteratively, with the bounds for each iteration shown as colored vertical lines.

4.4.4 Plane-to-plane correction

The truncated mean is used for the channel-to-channel correction because its low requirement of statistics. However, the ideal correction would be to equalize the peak muon ionization energy in each channel, since this quantity is roughly independent of muon energy. In a study integrating over long time periods and many strips, the truncated mean was found to be 4% higher than the peak value on average. The truncated mean is a good proxy for the peak if this 4% difference were consistent across the entire detector since the normalization of the strip-to-strip constants is

fixed.

Differences in the relationship between the truncated mean and the peak could occur if the shape of the energy distribution is not consistent between channels. Scintillator aging and absorber effects are expected to be among the leading causes of such shape differences. Scintillator within a plane was manufactured and installed at the same time, so aging effects might exist between planes but should be uniform within a plane. In the downstream calorimeters, absorber effects are uniform across a plane. In the nuclear target region, there are non-uniform absorbers that could introduce differences in the propagation of secondary electrons within a plane. Also, the side electromagnetic calorimeter could alter the distribution for strips which are contained entirely in that region. These effects are expected to be small, so that correcting to the peak plane-by-plane is sufficient.

By aggregating over the 127 scintillator strips within a plane, sufficient statistics are accumulated to perform a fit for the peak energy. Separately, the truncated mean procedure described in Section 4.4.3 is performed to extract the truncated mean energy for the entire plane. The correction factor C_i for plane i is

$$C_i = \frac{\frac{x_i}{p_i}}{\frac{1}{n} \sum_k \frac{x_k}{p_k}} \quad (4.5)$$

where x_i is the truncated mean energy averaged over a plane, p_i is the fitted peak energy for a plane, n is the number of planes and the sum in the denominator is over planes. The product of Equations 4.4 and 4.5 gives the applied strip-to-strip constant.

The fitted peak is determined by first fitting the energy per path length distribution to a quadratic polynomial

$$c_0 + c_2(x - c_1)^2 \quad (4.6)$$

where c_0 , c_1 and c_2 are free parameters in the fit and x is the energy per unit path length. The energy per path x is fit between the two values which correspond to half the maximum number of entries, and the parameter c_1 is constrained to take values in that range. The results of the quadratic fit are used to seed parameters in a fifth-order polynomial fit to the same distribution:

$$c_0 + c_2(x - c_1)^2 + c_3(x - c_1)^3 + c_4(x - c_1)^4 + c_5(x - c_1)^5 \quad (4.7)$$

where again the c_n s are free parameters. The seed value for c_1 and c_2 are taken from the fitted value of the fit to Equation 4.6. An example of a very similar fit in the absolute energy calibration is shown in Figure 4.19.

The calibration is iterative because applying a correction to the energy in each strip affects the relative energy distribution of adjacent hits in the same plane. This in turn affects the fitted position of a track in each plane, which affects the path length in each strip. The calibration is improved by performing a second iteration to account for these effects. To save processing time, except at the beginning of a run period, the constants from the previous interval of validity are used as a first iteration. Changes in the constants are small from one interval to the next, except when hardware is replaced, so the previous interval is an excellent starting point. Figure 4.14 shows the truncated mean energy per unit path length in units of MeV/cm initially and after one iteration. Figure 4.15 shows the final calibrated peak energy for each plane.




Figure 4.14: The uncalibrated data has a width of $\sigma = 19.1\%$, which decreases to $\sigma = 1.6\%$ after the first iteration. The second iteration reduces the width to $\sigma = 0.5\%$.

4.5 Absolute energy calibration

The absolute energy scale calibration ensures that the energy and photon statistics are the same in data and simulation. The calibration sample consists of through-going rock muon tracks which match to tracks in the MINOS detector with good charge and momentum reconstruction. A simulated sample is generated from the reconstructed muon four-vectors observed in data, such that the energy spectrum of muons is identical in data and simulation by construction. The calibration measures the relationship between photoelectrons and deposited energy, and two important




Figure 4.15: The peak energy per unit path length is fitted for each plane. The resulting peaks are fitted to a linear distribution with zero slope since the plane-to-plane energy is expected to be flat. The p-value for this fit is 0.90.

quantities are extracted. First, the “light yield” (LY) factor which governs the simulated photostatistics, and second, the muon energy unit or “MEU” factor, which governs the energy scale in data and simulation.

The calibration uses the peak “cluster” energy due to muons. A cluster is a set of adjacent hits in one scintillator plane, and a muon cluster is typically two hits wide. The function used to determine the peak is the same as described in the text leading to Equation 4.7. Muons are used because they deposit relatively little energy as they traverse the detector, and this energy deposit is only weakly dependent on




Figure 4.16: The peak cluster PE for rock muons in data decreases as a function of time due to scintillator degradation. The gaps represent shutdowns where no beam data was taken.

the energy of the muon itself.

The LY factor is the number of photoelectrons per MeV of muon energy deposited in the MINERvA scintillator. It is a combination of the number of photons per MeV from the scintillator, the absorption and re-emission by the fiber, and the quantum efficiency of the PMT. This factor is determined by generating a simulated sample with a trial light yield factor, and tuning the trial factor by the ratio of the peak cluster PE in data to simulation. This is done in roughly 2-day intervals to correct for a decrease in the light level with time due to scintillator degradation. This can




Figure 4.17: After applying the absolute energy calibration, the cluster energy distributions agree in data and simulation.

be seen in the peak cluster PE for data, shown in Figure 4.16.

The MEU factor is the energy per photoelectron. It is determined so that the peak cluster energy in data and simulation are equal to each other, and equal to the true simulated energy. The shape of the true energy distribution can be seen in Figure 4.18, and is very narrow compared to the reconstructed energy due to detector resolution. For this reason, the true energy is not fit directly to determine the peak. Instead, the simulated reconstructed energy distribution is fit to determine the peak, and a correction is computed by fitting the reconstructed versus true energy to a straight line. An example of the fit to reconstructed energy is shown in Figure 4.19.




Figure 4.18: The true cluster energy distribution in simulation. The peak due to minimum ionization is much narrower than the reconstructed distribution. The tail is due to processes other than ionization.

An example of the linear fit of reconstructed and true energy is shown in Figure 4.20. The fit is performed between 2 and 20 GeV only. As can be seen in Figure 4.18, clusters with true energy less than 2 GeV are rare. There is a reconstruction bias toward overestimating the cluster energy for these very low energy clusters because clusters below 1 MeV are not used in the tracking algorithm and can be missed.

The cluster energy distributions in data and simulation after the calibration are shown in Figure 4.17. The MEU factor is defined as




Figure 4.19: An example of the fit used in the strip-to-strip and absolute energy calibrations, shown here for reconstructed cluster energy in the absolute energy calibration.

$$\text{MEU} = \text{MEU}_{\text{trial}} \frac{E_{\text{reco}}^{\text{MC}}}{E_{\text{reco}}^{\text{data}}} \frac{E_{\text{true}}^{\text{MC}}}{E_{\text{reco}}^{\text{MC}}} \quad (4.8)$$

where $\text{MEU}_{\text{trial}}$ is the trial MEU factor, E_{reco} is the peak reconstructed cluster energy in data or Monte Carlo simulation, and E_{true} is the true simulated energy. The first factor comes from fits to the reconstructed energy distributions in Figure 4.17, while the second factor is the slope in Figure 4.20.

The absolute energy calibration corrects the data and simulated energy scales based on the energy deposition of a muon. It does not correct for the detector energy




Figure 4.20: Reconstructed cluster energy in simulation as a function of true cluster energy. The resulting linear fit is used as a correction.

response to other particles, nor does it account for energy due to neutral particles that are not observed. The energy response due to pions, protons, and electrons is studied in a miniature version of the MINERvA detector exposed to a beam of hadrons of known momenta [96]. The test beam detector energy scale is set by a procedure functionally identical to that used in the main detector, such that the observed energy due to other particles can be directly applied.

4.6 Timing calibration

The timing calibration corrects two effects. The first is a static timing offset for each FEB, and the second is a pulse-height dependent effect called “time slewing.” Each of MINERvA’s 507 FEBs has its own internal clock. Offsets between FEBs are typically of order 10 ns, and must be corrected. These offsets do not drift appreciably over time, but can change whenever the CRIM is power-cycled.

Each photon that arrives at the PMT photocathode is the product of two decays. First, the scintillator fluor decays, emitting a photon which is then absorbed by the fiber. Second, the fiber re-emits a photon, which is internally reflected inside the fiber. The photon either traveled directly from its point of origin along the fiber to the PMT, or it went in the opposite direction, and was reflected at the mirror before traveling the entire length of the strip to the PMT. For low pulse-height hits of just a few photoelectrons, these effects are significant. At higher number of PE, the probability that the first PE is due to direct rather than reflected light approaches unity, and the average decay time for the first PE goes to zero. Because of this, the timing resolution is 10 ns for 1-2 PE hits, 3 ns for 6-12 PE hits due to minimum ionizing particles, and approaches the 2.2 ns resolution of the electronics asymptotically at very high pulse heights. Example timing distributions for different pulse heights are shown in Figure 4.21.

In order to disentangle the electronics and the time slewing effects, the timing calibration is iterative. In each iteration, the corrections from the previous iteration are applied. Both outputs converge after several iterations, and 7 are used.




Figure 4.21: Area normalized timing histograms for four PE bins corrected for everything except time slewing.

4.6.1 The reference track time

The timing calibration sample consists of front-entering, through-going rock muons. For each muon track, a reference time is computed. While the uncertainty on the reconstructed time for a single hit is several ns, a much more precise time can be computed by averaging over many hits along a track. The relative time between two hits is known because the muon velocity is within 1% of the speed of light. Hit times are also corrected for the light time-of-flight in the optical fiber, with an assumed speed of 156 mm/ns based on a measured index of refraction of 1.923. The exact point of the muon passing through the strip is known from the track fit, so the fiber

path length is known precisely.

The reference time is the iterative, weighted truncated mean time of all hits on the track. In each iteration, hits are considered if they are within 1 standard deviation of the mean hit time from the previous iteration. In the first iteration, hits are weighted by PE. In subsequent iterations, the weight is the inverse of the standard deviation of the hits in that PE bin from the previous iteration. Hits below 3 PE are not considered.

4.6.2 Electronics offsets

The electronics offsets are computed for each block of 16 pixels, corresponding to one TriP-t chip. Each FEB has four such blocks, and variation among TriP-t chips on one FEB is typically less than 1 ns. The offset is calculated by forming a histogram of the difference between the hit time and reference track time for each hit in a 16-pixel block. At least 25,000 muons are used to create one calibration table, and since in general each muon will deposit energy in two strips in each plane, approximately 12,500 hits in each block are used.

The offset is the iterative truncated mean of the histogram, computed in the same way as the reference track time. For filling this histogram, hit times are corrected for everything except the electronics offset. This includes time-of-flight of the muon, time-of-flight of the light propagating in the fiber, and the time slewing correction described in Section 4.6.3. In the first iteration of the timing calibration (not to be confused with the iterative process by which the truncated mean is determined), there is no time slewing correction to apply. In later iterations, the time slewing effect is corrected, and the histogram for the electronics offsets becomes narrower. This can be seen in Figure 4.22, which shows 4 examples of how the histogram for




Figure 4.22: Electronics offsets at various point in the iterative timing calibration for the first 16 pixels of board 0-1-0-2.

one block of 16 pixels progresses. This board appears to be offset by about one clock tick (9.48 ns) early relative to the average over all boards in the detector. The reference track time is computed using all hits, and therefore the average time of all hits over any particular board is not necessarily equal to the reference time.

4.6.3 Time slewing correction

The time slewing correction is applied to hits based on their number of measured PE, and is the same for the entire detector. Hits are corrected for time-of-flight, and the electronics offsets are applied as measured in the previous iteration of the

calibration. Hits are put into bins of PE that are 0.5 PE wide below 2 PE, 1 PE wide from 2-30 PE, 2 PE wide from 30-40 PE, and 5 PE wide from 40-50 PE. Hits above 50 PE are rare for rock muons and are placed into an overflow bin. In each bin of PE, histograms are filled with the difference between the hit time and the track reference time. A fit is performed using a convolution of a Gaussian and a Landau distribution [97]. The Landau distribution accounts for the non-Gaussian high-side tail. The peak of the Landau-Gauss fit is extracted.

The distribution becomes narrower and more Gaussian at higher pulse heights, as the electronics resolution begins to approach the mean decay time of the scintillator and fiber. It also shifts closer to zero. The time slewing correction is defined to be zero at infinite PE. This definition is completely arbitrary, as it will be absorbed into the board offsets in subsequent iterations of the calibration. Timing distributions corrected for everything except time slewing are shown in Figure 4.21 for four example PE bins.

The Landau-Gauss peak is obtained for each bin. The peak is then fit as a function of PE with a third-order polynomial in $1/\sqrt{PE}$. The result of the fit can be seen in Figure 4.23. The fit considers only the range from 2 to 50 PE to avoid undefined behavior at 0 PE. Because of the pedestal subtraction threshold, no hit is ever measured to have a pulse height below about 0.3 PE. For analysis, timing of hits below 3 PE are not considered. The fit parameters are stored, and the time slewing correction that goes into the calibrated hit times comes from the fit.

The hit time residual as a function of pulse height with all timing corrections applied is shown in Figure 5.7.




Figure 4.23: The Landau-Gauss peak time slewing in each bin of PE is fit to determine the time slewing correction.

Chapter 5

Simulation

Event simulation in MINERvA consists of three parts. First, neutrino interactions with the nuclei in the detector are simulated. Second, particles exiting the nucleus are then passed to a detector simulation package, which propagates the particles through the detector. Third, custom simulation software is used to mimic the read-out of the detector and produce calibrated, time-stamped hits similar to what is seen in data. Once this is done, identical reconstruction and analysis software is run on data and simulation.

5.1 Neutrino interaction simulation

Neutrino interactions in the MINERvA detector are simulated with the GENIE (Generates Events for Neutrino Interaction Experiments) version 2.8.4 package [59]. A specified number of protons on target are simulated. GENIE randomly selects neutrinos based on the flux simulation described in Section 3.2. A geometry description of the MINERvA detector is constructed. This geometry includes the scintillator

planes, the lead collar, the upstream nuclear targets, the liquid helium target and cryostat, and even the air gaps inside the detector. Because MINERvA analyses select events either from the tracker or the upstream nuclear targets, interactions in the calorimeters are not simulated. Backgrounds from events originating in the outer detector steel can be rejected by requiring that there not be any energy in the outer detector. Due to the very large mass of the calorimeters as compared to the tracker and nuclear targets (the hadronic calorimeter has a total mass of 146 tons, compared to 8.3 tons for the tracker), simulating events proportionally in these subdetectors would be computationally expensive.

A ray is drawn in the neutrino direction, and an interaction point is chosen along the ray at random, weighted by the total interaction cross section in the materials intersected by the ray. Once the target atom is selected, an interaction type is chosen randomly. By default, GENIE models quasi-elastic scattering, resonance production, deep inelastic scattering, and coherent pion production, for both charged and neutral currents. It also includes three rare processes: charged-current charm production, neutrino-electron elastic scattering, and inverse muon decay $\nu_\mu e^- \rightarrow \nu_e \mu^-$.

The nuclear model used in GENIE is the relativistic Fermi gas (RFG) [52]. It assumes nucleons inside the nucleus are quasi-free, acting independently in the mean field of the nucleus. The initial-state momentum of the struck nucleon inside a nucleus comes from the Fermi distribution, modified with a high-momentum tail to account for correlations between nucleons using the Bodek-Ritchie model [53, 54]. Quasi-elastic scattering is predicted using the Llewellyn Smith model [50], which parameterizes the cross section in terms of nucleon form factors. Baryon resonance production for hadronic invariant mass $W < 1.7$ GeV comes from the Rein-Sehgal model [98]. Deep inelastic scattering is simulated with the Bodek-Yang model [99].

For each process, the total cross sections at various neutrino energies are stored in spline files. The splines are interpolated to form continuous functions, which are called for the neutrino energy selected from the flux simulation in order to select an event type. Once the event type is selected, the kinematics are generated. The particles produced in the interaction are placed inside the target nucleus. They are then passed to a module called INTRANUKE [59, 100]. This package simulates final-state interactions (FSI), where hadrons produced in the neutrino interaction subsequently interact while inside the struck nucleus. Some other event generators [25, 58] use a full intranuclear cascade model to simulate FSI. The model used by MINERvA is the GENIE “hA” model, an effective cascade model, which represents the effect of the full cascade in a single interaction. Hadrons born inside the nucleus in GENIE undergo either zero or one interaction. The INTRANUKE model is tuned to external hadron-nucleus scattering data.

5.1.1 Simulation of kaon production

A summary of K^+ production in GENIE is given in Table 5.1. Kaons are produced via the decays of baryon resonances as well as from hadronization in DIS events. In GENIE, individual resonances are simulated only for hadronic invariant masses $W < 1.7$ GeV, which is just above the $K^+\Lambda$ threshold. Most K^+ originate in hadronization, the process of creating hadrons from the struck quark. For $1.7 < W < 2.3$ GeV, the hadronization model is a parameterization based on Koba-Nielsen-Olesen (KNO) scaling [101]. PYTHIA6 [102] is used for $W > 3.0$ GeV. In the intermediate region $2.3 < W < 3.0$ GeV, the AGKY model [103] governs the transition between KNO and PYTHIA6. Parameters which control the rate of strange particle production in hadronization are tuned such that rates of Λ and K_S^0 pro-

$\nu_\mu N \rightarrow \mu^- K^+ N$	In GENIE version 2.10 and later
$\nu_\mu N \rightarrow \mu^- \Delta \rightarrow \mu^- K^+ Y$	In GENIE for $W < 1.7$ GeV
$\nu N \rightarrow \nu \Delta \rightarrow \nu K^+ Y$	
$\nu_\mu q \rightarrow \text{hadronization} \rightarrow \mu^- K^+ Y$	KNO ($1.7 < W < 2.3$ GeV) AGKY ($2.3 < W < 3.0$ GeV) PYTHIA6 ($W > 3.0$ GeV)
$\nu q \rightarrow \text{hadronization} \rightarrow \nu K^+ Y$	
$\nu_\mu q \rightarrow \text{hadronization} \rightarrow \mu^- K^+ \bar{K}$	
$\nu q \rightarrow \text{hadronization} \rightarrow \nu K^+ \bar{K}$	
$\nu N \rightarrow \text{non-strange} \xrightarrow{FSI} K^+ X$	not simulated

Table 5.1: A summary of K^+ production reactions simulated in GENIE. Muon-flavor CC reactions are listed, but the analogous electron-flavor process also occurs. In the table, N is a nucleon, Y is a hyperon, and \bar{K} is a K^- or \bar{K}^0 .

duction on deuterium and nuclei agree with BEBC [60, 61, 62, 63] and Fermilab 15' [64, 65] bubble chamber measurements as a function of W .

In principle, the K^+ rate in GENIE could be modified by changing the total DIS rate. However, the DIS cross section, unlike the cross section for K^+ production in particular, is well constrained by data at high neutrino energy. There are three parameters in GENIE which control the rate of K^+ production relative to other particles in DIS. In the KNO and AGKY portions of the hadronization model, the probability to produce a strange baryon is

$$0.021951447 + 0.041969985 \log W, \quad (5.1)$$

and is cut off at 1.0 arbitrarily, which is entirely irrelevant since the cutoff corresponds to a hadronic invariant mass of 10^{11} GeV. These parameters come from fits to data with 10-30% uncertainties, and are not really known to a few parts in 10^8 as the precision would suggest. Because these parameters control Λ and Σ baryons in associated production, they effectively also control the strange meson (either a K^+ or K^0) produced along with the baryon. There are also parameters that control

meson pair production. When the available phase space allows it, $\pi^+\pi^-$ pairs are produced 63% of the time, with $\pi^0\pi^0$ at 31%, and K^+K^- and $K^0\bar{K}^0$ at 3% each. At higher W , there is a parameter called s which controls the suppression of $s\bar{s}$ pairs relative to $d\bar{d}$ pairs being pulled from the vacuum in PYTHIA. It is set to 0.30 in GENIE based on the bubble chamber data at high W . In NUX [104], the generator used in the NOMAD experiment which also measured K^0 and Λ production, PYTHIA was also used for high- W hadronization but with s set to 0.21.

Kaon rescattering was added to the INTRANUKE FSI model in version 2.8.0, and is tuned to data from Bugg *et al.* [105] and Friedman *et al.* [106]. Charge exchange processes are not included for kaons, but are part of the model for pions. K^+ absorption is forbidden due to strangeness conservation because it has strangeness $S = 1$ and strange baryons are all $S = -1$. Processes such as $\pi^+n \rightarrow K^+\Lambda$, in which a K^+ is produced in association with another strange particle, are not simulated in GENIE. The only effect of FSI on K^+ production in GENIE is to decrease the K^+ energy when it rescatters. The number of K^+ in the final state is never modified by FSI in GENIE 2.8.4.

Figure 5.1 shows the total cross sections in GENIE 2.8.4 for K^+ production with K^+ kinetic energy less than 600 MeV. The cross sections are broken up into resonance production and DIS, with DIS further subdivided into KNO, AGKY, and PYTHIA components. For the flux of MINERvA, the number of events from KNO, AGKY, and PYTHIA are approximately equal.

5.1.2 Single kaon production in genie

GENIE 2.8.4 has no model for strangeness-nonconserving processes. Single kaon production, $\nu_\mu N \rightarrow \mu^- K^+ N$, is an important contribution to the total K^+ cross




Figure 5.1: Charged-current (left) and neutral-current (right) total K^+ cross sections as a function of neutrino energy for $T_K < 600$ MeV only, broken up based on which model was used to generate the event. The NC plot includes the antineutrino flux and combines neutrino and antineutrino events.

section at low neutrino energies (1-3 GeV). With no final-state hyperon, this process has a threshold of 0.8 GeV for free nucleons, compared to 1.2 GeV for associated K^+ production. Near threshold, this process could contribute to the background for $p \rightarrow K^+ \bar{\nu}$ if the μ^- is below Cherenkov threshold and captures.

Single kaon production is predicted by Alam *et al.* [107]. The model includes four diagrams, shown in Figure 5.2. The contact term (top-left) is largest in this model, and is actually larger than the total cross section but interferes destructively with the other diagrams. The kaon pole term (top-right) is essentially a W boson fluctuating to a kaon, which is subsequently brought on-shell by an interaction with a nucleon. This term is suppressed by the lepton mass and contributes negligibly to the total cross section. In the u-channel cross baryon diagram (bottom-left), the reaction proceeds by an intermediate-state Λ or Σ baryon. There is no s-channel with such an intermediate state due to the lack of baryons with $S = 1$ (in the analogous antineutrino reaction $\bar{\nu}_\mu N \rightarrow \mu^+ K^- N$, this diagram does contribute).




Figure 5.2: Feynman diagrams for single kaon production considered in the model described in Ref. [107]. In reading order: the contact diagram, kaon pole diagram, cross baryon diagram, and meson in-flight diagram.

Finally, the meson in-flight term (bottom-right) involves the exchange of a pion or eta meson.

While the figures in Ref. [107] show only one-dimensional cross section predictions, the model is in four dimensions. It gives the four-vectors of the outgoing lepton, kaon, and nucleon. Four-momentum conservation and rotational invariance about the neutrino axis reduce the dimensionality from nine to four. It never produces additional final-state hadrons, which is not realistic at higher W . As the neutrino energy increases, it is expected that additional pions would be added to the final state, for example $\nu_\mu N \rightarrow \mu^- K^+ \pi^0 N$ or $\nu_\mu N \rightarrow \mu^- K^+ \pi^+ \pi^- N$. The calculations in Ref. [107] are valid at low neutrino energy, and the range of validity suggested by the authors is up to 2 GeV. At high neutrino energies, the chiral perturbation theory used in the model breaks down.

This prediction is implemented in GENIE beginning with version 2.10.0 [108], for neutrinos only (not for antineutrinos), as an optional model. In version 2.10, the

suggested range of validity is ignored, and the model will generate valid events at any neutrino energy, including highly unlikely two-hadron final states at very high W .

For the MINERvA analyses, an alternate simulated sample is constructed by including the $\Delta S = 1$ events generating using GENIE 2.10.0, and weighting down the $\Delta S = 0$ events generated using GENIE 2.8.4 in order to preserve the total cross section as a function of W . Strange particle production in GENIE is tuned to data as a function of W , and this procedure attempts to preserve that tuning. An empirical fit to the simulated cross sections gives

$$\frac{\sigma_{\Delta S=1}(W)}{\sigma_{\Delta S=0}(W)} \approx -0.016 + \frac{0.28}{W^{1.09}}. \quad (5.2)$$

The weight applied to all $\Delta S = 0$ events with $W < 4$ GeV is given by

$$\text{wgt} = 1 - \frac{\sigma_{\Delta S=1}(W)}{\sigma_{\Delta S=0}(W)}. \quad (5.3)$$

Events with $\Delta S = 0$ are weighted down by an average of 7.7% to preserve the total K^+ production cross section when the $\Delta S = 1$ events are added. A comparison of the cross sections for $\Delta S = 0$ and $\Delta S = 1$ K^+ production as a function of K^+ kinetic energy is shown in Figure 5.3.

5.2 Detector simulation

The detector simulation begins with a Geant4-based program to simulate energy deposits in the active material of MINERvA. Then, a custom program converts those energy deposits into recorded hits similar to those observed in data.




Figure 5.3: The cross section for single kaon production is shown as a function of K^+ kinetic energy compared to the total associated production cross section.

5.2.1 Geant4 simulation

The final output of the GENIE stage of the simulation is a list of particle four-vectors as they exit the nucleus. This is passed to the Geant4 (GEometry AND Tracking) program [74, 75], which steps particles through the detector medium. Version 9.4.p02 is used in MINERvA. For the Geant stage, the full detector geometry is used, including the downstream calorimeters and outer detector. In addition, a distinction is made between active material (scintillator), where energy deposits will be recorded, and passive materials, where they will not be.

In the MINERvA simulation, Geant steps particles 1 mm at a time. At each




Figure 5.4: The K^+ -carbon inelastic cross section in Geant4 version 9.4.p02 does not agree with external data at low energy.

step, a decision is made randomly whether to have the particle interact, decay, or simply continue to propagate, depositing ionization energy if it is charged. These decisions are made based on cross sections compiled in a “physics list.” MINERvA uses the QGSP_BERT list, which combines the quark gluon string precompound (QGSP) model [109] with a Bertini cascade [80, 81]. For particles produced in neutrino interactions given the MINERvA flux, the QGSP portion of the model is rarely used. It covers pion, proton, neutron, and kaon interactions above about 10 GeV/c of momentum. The Bertini cascade is used below 10 GeV/c, and like GENIE’s INTRANUKE model it is tuned to hadron-nucleus scattering data where available.

The K^+ -carbon inelastic cross section in the Geant4 version used by MINERvA has an unphysical wiggle at low K^+ kinetic energy and is in disagreement with external data [105, 106]. Figure 5.4 shows the elastic and inelastic components of the cross section. Between 200 and 400 MeV, the data rises slowly while the Geant4 prediction increases rapidly before falling back into agreement with total cross section data between 400 and 800 MeV. Kaon interactions in the MINERvA detector affect the measured K^+ kinetic energy and are very important for the measurement of the K^+ spectrum.

To correct for this feature, a weight is applied to simulated events based on whether the K^+ scatters inelastically, elastically, or not at all, according to Geant4. The weights are given in Eq. 5.4:

$$\begin{aligned}
 W_{inel} &= \frac{1 - e^{-\rho x \sigma_{data}^{tot}}}{1 - e^{-\rho x \sigma_{geant}^{tot}}} \times \frac{\sigma_{data}^{inel}}{\sigma_{data}^{tot}} \times \frac{\sigma_{geant}^{tot}}{\sigma_{geant}^{inel}} \\
 W_{el} &= \frac{1 - e^{-\rho x \sigma_{data}^{tot}}}{1 - e^{-\rho x \sigma_{geant}^{tot}}} \times \frac{\sigma_{data}^{el}}{\sigma_{data}^{tot}} \times \frac{\sigma_{geant}^{tot}}{\sigma_{geant}^{el}} \\
 W_{none} &= e^{-\rho x (\sigma_{data}^{tot} - \sigma_{geant}^{tot})},
 \end{aligned} \tag{5.4}$$

where ρ is the density of the tracker and x is the distance traveled by the K^+ . The Geant4 prediction σ_{geant}^{el} (σ_{geant}^{inel}) is taken from a spline fit to cross sections determined by counting elastic (inelastic) interactions in a simulation of K^+ incident on a thin carbon target. The total cross section prediction σ_{geant}^{tot} is the sum of the elastic and inelastic components. The data constraint σ_{data}^{tot} is a parameterization of K^+ -carbon total cross section measurements [105, 106]. The inelastic constraint σ_{data}^{inel} is a parameterization of reaction cross section data [106], and includes nucleon knock-out. The elastic component σ_{data}^{el} is not measured directly. Its shape as a

function of K^+ energy is taken from Geant4, and it is normalized to agree with the average difference between the total and inelastic data. Scattering on other nuclei in the tracker is reweighted based on the carbon data and A -dependent nuclear effects are not considered.

For events with K^+ kinetic energy less than 600 MeV, 16% undergo only elastic scattering (reduced to 10% by reweighting), 28% experience inelastic reactions (increased to 34%), and the remaining 56% of events have no K^+ interaction. After reweighting, the inelastic and total scattering cross sections as a function of kaon energy agree with external K^+ -carbon scattering data, as shown in Fig. 5.5. The wiggle is not present in the newest version of Geant4 10.0 [110].

5.2.2 Data overlay

Simulated events in MINERvA are overlaid with data gates prior to reconstruction. This is done primarily to simulate pile-up, where two neutrino interactions in the same beam spill overlap in space and time to fool the reconstruction, or to cause an event to not be reconstructed. The data overlay occurs prior to the creation of time slices, described in Section 4.3. The effect of pile-up on the K^+ event selection is discussed in Section 6.2.

A second effect of the data overlay is the implementation of time-dependent calibration constants in the simulation. Data is selected at random, weighted by the number of protons on target. The calibrations used in the data gate that is overlaid on a simulated interaction is also used in the simulation. Two important time-dependent effects are the light level and dead channels. The light level decreases with time, reducing the number of observed photoelectrons and increasing the statistical fluctuations in the number of PE. Dead channels encountered in data are masked




Figure 5.5: Kaon interactions predicted by Geant4 are reweighted to agree with external data from Bugg *et al.* [105] and Friedman *et al.* [106] on the total and inelastic scattering cross sections on carbon.

in the simulation, and energy deposited in those strips by the Geant simulation is ignored.

5.2.3 Readout simulation

The goal of the readout simulation is to make the simulated quantities as realistic as possible. It begins with true energy deposits in active materials from Geant4, and “decalibrates” those energy deposits, essentially inverting the calibration procedure described in Chapter 4, to arrive at raw ADC counts. There, the pedestal suppression and discriminator thresholds are applied, as they are in data. Finally,

the raw counts are calibrated by the same procedure as in the data to arrive at the reconstructed simulated energy.

5.2.4 Optical model

Light is simulated using Birks' law [111], which governs the relationship between ionization energy and light emitted by a scintillator:

$$\frac{dL}{dx} = L_0 \frac{\frac{dE}{dx}}{1 + k_B \frac{dE}{dx}}, \quad (5.5)$$

where L is the amount of light, dE/dX is the energy deposited per unit distance, and k_B is Birks' constant. The base level L_0 is the light observed when dE/dx is small and the light is proportional only to the energy deposited. The value of L_0 used in the MINERvA simulation is 8000, though regardless of the chosen value of L_0 , the light yield factor determined from data in the absolute energy scale calibration will give the number of photoelectrons. The Birks' constant k_B is found to be 0.0905 ± 0.0150 by analyzing a sample of stopping protons in the MINERvA test beam detector, which uses the same scintillator as MINERvA [96].

The Birks-suppressed light is converted to a mean number of photoelectrons using four factors. First, the number of photons is divided by the strip-to-strip factor. Second, the number of photons is multiplied by the total attenuation. The green fiber "baggie" attenuation and clear fiber attenuation depend only on the strip number. The attenuation in the scintillator strip is applied based on the true location of the energy deposit. Half of the light is assumed to travel directly to the readout end of the strip, while the other half travels to the mirror then reflects before traveling the length of the strip. Third, the photons are converted to PE by

multiplying by the light yield factor.

The fourth and final factor is a channel-to-channel smearing factor. This is a deliberate smearing applied to bring the width of the simulated energy distribution for rock muon clusters into agreement with data. The channel scale factors are determined randomly from a Gaussian with $\mu = 1.0$ and $\sigma = 0.0557$. The value of σ is determined empirically by generating simulated samples with Gaussian smearing of different widths between 0.0 and 0.1 in steps of 0.005. The χ^2 between data and each simulated sample is calculated, and the optimal channel-to-channel smearing is determined from a fit to be 0.0557. The smearing represents the total uncertainty in the entire calibration procedure. The uncertainty in the strip-to-strip constant for a single strip is estimated to be 1.5%. There is also an observed statistical fluctuation in the measured gain for a single pixel of 3-5% that decreases when more light injection data is used. The 0.0557 comes from a combination of these, as well as possibly other sources.

The actual number of photoelectrons is then chosen randomly from a Poisson distribution, with mean equal to the Birks-suppressed number of photons multiplied by the four factors described above. Timing is not generated from a photon by photon model, but instead from a parameterization described in Section 5.2.6.

5.2.5 PMT and ADC model

The photoelectrons output by the optical model are then simulated in the photomultiplier tube in a PMT model. The primary purpose of this stage of the simulation is to add cross-talk. The signal is also amplified by the channel gain as measured in data at the time of the overlaid data gate.

Cross-talk is when a signal from one channel leaks onto another. It is dominated

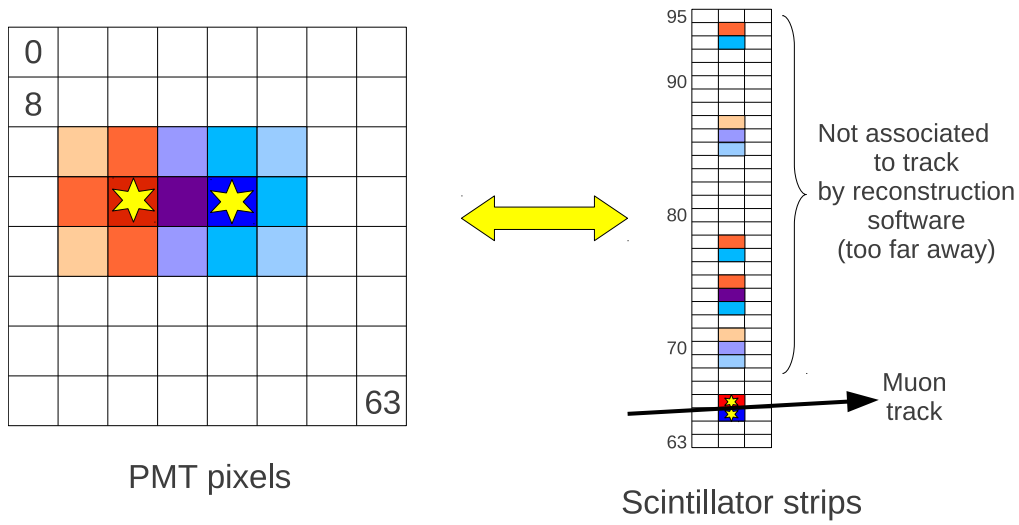


Figure 5.6: Optical cross talk is weaved from neighboring channels on the PMT to strips in the detector that are not adjacent.

by “optical cross-talk,” when light from a fiber that is meant to be coupled to a channel results in photoelectrons on a neighboring channel on the PMT face. Channels that are neighbors on the PMT are “weaved” so that the scintillator strips they instrument are not adjacent, making it possible to algorithmically suppress cross-talk hits. An example of this is shown in Figure 5.6. A smaller effect is “dynode chain cross-talk,” where secondary electrons from one dynode are ejected at an angle and cause a cascade in a neighboring channel.

The amount of cross-talk differs between different PMTs. For example, if the PMT photocathode is misaligned with the optical fibers in a particular PMT, higher cross-talk would be expected. The cross-talk level is measured in situ from a sample of rock muons with no other activity in the time slice. Small energy deposits not on the muon track are assumed to be cross-talk. A source channel is identified by looking at hits that are reconstructed along the muon track. The total attenuation-corrected PE in strips which correspond to pixels adjacent to the source is summed,

and the cross-talk fraction is obtained by dividing the total cross-talk PE by the source PE. A scaling factor is determined for each PMT in data and applied to the simulation. After signals have been amplified and cross-talk hits added, a random amount of electronics noise is added to the event. Cross talk simulation and identification is described in detail in Ref. [87].

The ADC model digitizes charges from the PMT using the three gain ranges, which can be seen in Figure 4.5. A pedestal “sparsification threshold” cut is applied, and only hits which exceed the pedestal mean by at least three standard deviations are kept. This imposes a minimum PE threshold of approximately 0.3, similar to what is seen in data. The hit times are also digitized into “ticks” and “quarter ticks” where the finest timing is in units of 2.2 ns. Hits are discarded if they occur during the dead time window of a previous hit. For simulated hits, it is extremely unlikely that a hit will be discarded due to dead time from a previous simulated hit because the time scale of one simulated event is much smaller than 151 ns. However, data overlay hits are also considered, and simulated hits can be removed when they occur during the dead window that is induced by data overlay.

Finally, the fully-decalibrated raw hits are passed back through the calibration chain described in Chapter 4.

5.2.6 Timing simulation

The timing simulation is data-driven. True hit times are digitized in the ADC model, and a smearing term is added based on the results of the timing calibration. A random smearing is selected for each hit directly from the histogram of the difference between the hit time and reference track time for that bin of PE. The two-dimensional histogram used to smear the simulated hit times is shown in Fig-




Figure 5.7: The histogram used to smear simulated hit times, derived directly from the timing calibration in data.

ure 5.7. It is the result of applying all time-of-flight corrections, the board offsets, and the time slewing correction described in Section 4.6, and represents the best guess at the true hit time. The residual is widest and most asymmetric at low pulse heights.

Chapter 6

Reconstruction and event selection

Reconstruction in MINERvA is divided into two stages. The first is the general reconstruction, which is independent of the specific neutrino interaction channel being studied. In this stage, long charged particle tracks are formed, and matching tracks in the MINOS near detector are identified. The neutrino interaction point, called the “primary vertex,” is also selected. The second stage is analysis-specific reconstruction. Results are reported for three different signal definitions:

1. Charged-current K^+ production: at least one K^+ exiting the nucleus, exactly one μ^- , and any other particles.
2. Neutral-current K^+ production: at least one K^+ exiting the nucleus, zero charged leptons, and any other particles.
3. Charged-current coherent K^+ production: exactly one K^+ and one μ^- , recoil nucleus intact.

In all three cases, the analysis-specific reconstruction begins by identifying a K^+ decay-at-rest candidate using timing information. After the K^+ has been selected,

events are separated into the three categories based on the presence or absence of a long track consistent with being a μ^- , and coherent events are identified by looking for the absence of nuclear break-up. The charged- and neutral-current analyses are qualitatively similar and will be described together in this chapter. The coherent analysis is very different and is described in Chapter 8.

6.1 General reconstruction

Hits in adjacent strips are grouped into clusters. These hits are typically due to a single charged particle, which traverses two strips as it passes through a scintillator plane because of the strips' triangular shape (see, for example, Figure. 4.10). Clusters are formed from one or more strips. For multi-hit clusters, the position is the energy-weighted centroid of the strips in the cluster. For muon tracks, the cluster position resolution is 3 mm, while the strip centers are 17 mm apart.

Clusters are then classified into five categories: low-activity, trackable, heavy ionizing, supercluster, and cross-talk candidate. Examples of energy deposition patterns and the resulting cluster classification are shown in Figure 6.1. Low-activity clusters are those with total energy deposits less than 1 MeV. A trackable cluster is what would be expected from a minimum ionizing particle. It must have energy between 1 and 12 MeV, with 1-3 adjacent strips at or above 0.5 MeV. Heavy ionizing clusters are similar to trackable clusters but have no maximum energy. These clusters are often found near the endpoint of hadron tracks, with large but narrow energy deposits. Superclusters are those with more than 1 MeV of energy that do not meet the criteria for trackable or heavy ionizing clusters. Typically this is because there are 4 or more strips with more than 0.5 MeV. Superclusters are common in

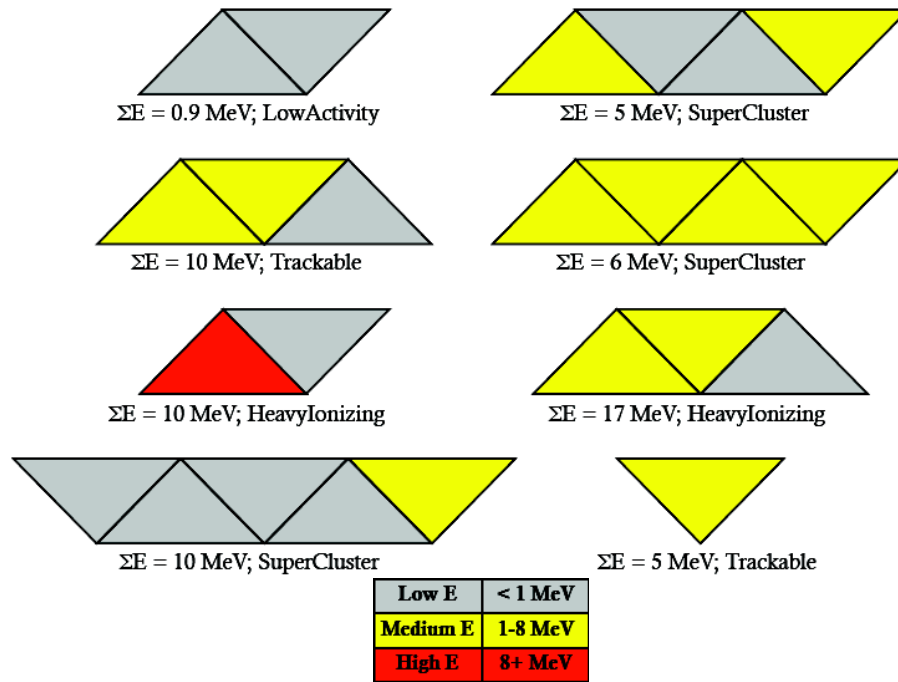


Figure 6.1: Examples of different energy deposition patterns, and the category to which the cluster is assigned in each case. Figure from G. Perdue.

developing electromagnetic showers. Finally, cross-talk candidates are identified by using the plex to project the pattern of energy deposits in a plane onto the PMT face, and looking at correlations between neighboring pixels. Low-energy hits in strips that are not adjacent to any other illuminated strip, but are adjacent to a large energy deposit on the PMT face, are typically due to cross-talk and are identified as such. There is no fixed energy range for cross-talk candidates but the hits are generally low energy. On average, 4% of a hit's charge will cross-talk to other pixels, so a cross-talk candidate cluster will be below 1 MeV unless the parent hit is larger than 25 MeV.

Track formation begins by identifying “seeds” where three clusters in consecutive planes of the same orientation form a two-dimensional line in either the XZ, UZ, or

VZ plane. Any three clusters can form a seed, so that the same cluster can be part of multiple track seeds. Then, seeds of the same view (X, U, or V) are merged if they have consistent slopes and intercepts to form two-dimensional tracks. The 2D track candidates are then stitched together to form three-dimensional objects. First a “three-view” track is formed if a 2D track in each of the three projections is found with the same longitudinal extent and consistent with a single three-dimensional line. After this is done, remaining 2D tracks are combined into “two-view” tracks if they are consistent with a 2D track in one other view. Additional clusters are then added in the third view. This second method is used when two tracks overlap in one view. All tracks are fit with a Kalman filtering algorithm that allows for multiple scattering [112]. Only trackable and heavy ionizing clusters are initially considered by the tracker. After the track is found, superclusters which intersect the track are added. The tracking algorithm is described in more detail in Ref. [4].

Separate from the MINERvA software, tracks are reconstructed in the MINOS near detector. Tracks exiting the rear of MINERvA are compared with tracks entering the front of MINOS. Matches are identified when the slope and intercept of a MINERvA and MINOS track match, and the tracks occur at the same time.

A MINOS-matched track is assumed to be a muon in the MINERvA reconstruction, and is treated as the “anchor” track in an event. The most upstream point on the anchor track is the initial guess for the neutrino interaction point. Other tracks are dissolved, and a second search is performed that is anchored to the muon vertex. If other tracks are found, the primary vertex is refit, and set to the point of closest approach of the tracks. Track matching residuals are shown in Figure 6.2. The timing residual is shown in Figure 6.3. Tracks are considered matched when the timing difference is within ± 50 ns.




Figure 6.2: Track matching resolutions (MINERvA track - MINOS track position) for the horizontal coordinate (left) and the vertical coordinate (right), projected to the rear of MINERvA (top) and to the front of MINOS (bottom). Figure from J. Chvojka.

6.2 Charged kaon reconstruction

Kaons are selected by reconstructing the timing signature of a K^+ decay-at-rest. This requires that the K^+ stop inside the tracker or electromagnetic calorimeter. If the K^+ stops in the hadronic calorimeter, 90% of the energy from its decay products is deposited in passive material and the K^+ cannot be reliably reconstructed. Non-interacting kaons with more than 600 MeV of kinetic energy typically reach the hadronic calorimeter and cannot be reconstructed using this timing-based technique. High-energy kaons are reconstructed only when they interact inelastically inside the tracker, in which case the range-based kinetic energy measurement is poor.

The timing signature reconstruction begins with a search for activity in the detector that is delayed in time with respect to the neutrino interaction, consistent




Figure 6.3: Track matching timing resolutions (MINERvA track - MINOS track time).

with the 12.4 ns K^+ lifetime, and consistent in energy with the products of a K^+ decay-at-rest (see Table 2.1). First, we search for a fully-reconstructed $K^+ \rightarrow \mu^+$ decay, which occurs in 64% of K^+ decays [5]. The tracking algorithm described in Section 6.1 is run anchored to the endpoint of all tracks that originate from the primary vertex. When secondary tracks are found, the event is classified as a “kinked track” event. An example of such an event is shown in Figure 6.4. If no kinked track is found, hits not associated with any tracks are grouped into narrow bunches in time, called “time slivers,” described in detail in Section 6.2.2. Events are accepted if there is a delayed time sliver that is spatially near the endpoint of a K^+ track. We also search for delayed time slivers near the neutrino interaction




Figure 6.4: A ν_μ -induced charged-current K^+ candidate in MINERvA data is viewed from above. The beam is angled into the page at 3.5° with respect to the horizontal axis. Each colored triangle represents one hit, a time-stamped energy deposit in a single scintillator strip. The color represents the hit time, relative to the reconstructed time of the interaction. The green circle is the event vertex, orange circles are reconstructed track endpoints, and the blue circle is a track kink. The green lines are reconstructed tracks. The kinked track is the K^+ candidate. The longest track is the muon candidate and is matched to a negatively-charged track in MINOS. The second segment of the kinked track is a μ^+ from the decay-at-rest $K^+ \rightarrow \mu^+ \nu_\mu$, with a time gap between the two segments of 18 ns. The remaining particles are likely the decay products of $\Sigma^+ \rightarrow \pi^+ n$, where the π^+ is the other track and the detached hits are proton products of a scattering neutron.

vertex. This extends the acceptance to K^+ with very small kinetic energy. These two samples are combined with additional selections to purify the K^+ content. The kaon selection methods are described in detail below.

6.2.1 Kinked track reconstruction

Topologically, kinked tracks due to $K^+ \rightarrow \mu$ decay are similar to pions which undergo a hard scatter in the detector and abruptly bend. Pions and protons vastly outnumber kaons in MINERvA, so the background is very large. Timing is the key

to distinguishing between these two cases. Hit times are corrected for time-of-flight and fit under two hypotheses. In the first hypothesis, the two segments are assumed to have the same true time, as would be the case for an interacting pion or proton. In the second hypothesis, the true times of the two segments are allowed to float independently. For signal events, the second segment is due to the μ^+ from K^+ decay, and will be late in time relative to the first segment, which is the K^+ track.

The probability density functions used in the timing fit are identical to those used in the simulation and examples are shown in Fig. 6.5. The fit maximizes the sum over all hits of the natural logarithm of the probability density. By construction, the two-parameter kaon decay hypothesis always gives a better fit. When the true time of the two segments is very small, the two hypotheses fit equally well, and the value of the log-likelihood is small. True decaying kaons will fit very poorly under the hypothesis where the two times are equal and result in a large log-likelihood score. Events are selected if the log-likelihood ratio is greater than 20.

An example signal candidate data event is shown in Fig. 6.4. The time gap distribution for a background-rich sample, and the log-likelihood ratio distribution, are shown in Fig. 6.6. The peak in the time gap plot is slightly below zero because the time-of-flight correction assumes a low-energy stopping kaon. Backgrounds with small time gaps are primarily due to hadron interactions in the detector. Pion interactions $\pi C \rightarrow \pi X$, where both segments of the kinked track are pions, will have small time gap and small log-likelihood ratio. Other interactions include proton knockout, $\pi C \rightarrow pX$, where the second segment is a proton. These events will also have small time gaps. The low-side tail in Figure 6.6a is due to events where the track direction is truly backward but is reconstructed as forward. In these events, the first and second segments are reversed, and the time-of-flight correction goes in




Figure 6.5: The probability density functions used in the simulation and for the timing fit are constructed using through-going muons in data. Three bins of PE are shown here as an illustration; in the actual fit, finer bins are used.

the wrong direction.

6.2.2 Time slivers

The slivering algorithm is very similar to the time slicing algorithm described in Section 4.3. Time slices are due to different neutrino interactions in the same 10 μs beam spill. Within each time slice, time slivers are formed. The goal of forming time slivers is so that a K^+ and its decay products are in two different slivers. When two neutrino interactions occur between ~ 10 and ~ 100 ns apart in time, they will be reconstructed in the same time slice but in different time slivers.

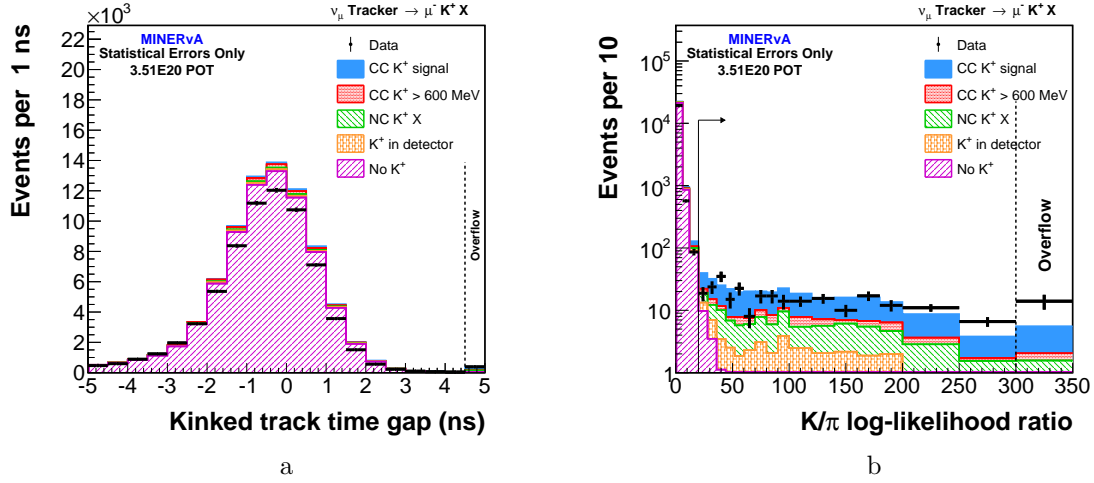


Figure 6.6: (a) The time gap between primary and secondary segments of a kinked track. The peak region is mostly due to interacting pions, where the deficit in data relative to simulation is consistent with other results indicating an overprediction in GENIE’s pion production model [113]. (b) The log-likelihood ratio of the fit described in the text separates stopping kaons from interacting hadrons. The arrow shows the selection of events with log-likelihood ratio greater than 20.

	Time slices	Time slivers
Window width	80 ns	5 ns
Hit metric	PE	weight defined in text
Threshold	10 PE	1.5 units
Minimum total	30 PE	1.5 units

Table 6.1: A summary of the differences between time slices and time slivers. The weight used in sliver formation is defined in the text.

In both algorithms, a fixed-width window slides forward in time. When the sum of all hits in the window exceeds a threshold, a slice or sliver is opened. The window continues to slide forward until the threshold is no longer met, and then the slice or sliver is closed. The differences between slice and sliver formation are summarized in Table 6.1.

The granularity of the search is defined by the width of the sliding window.

For the separation of different neutrino interactions in the same beam spill into time slices, this granularity is 80 ns. For the separation of a K^+ and its decay products into time slivers, a 5 ns window is used. In addition to the finer sliding window, a different metric is used to weight hits. For time slices, the total number of photoelectrons must exceed 10 in the 80-ns window for a slice to be opened, and a slice is kept only when the total exceeds 30 PE.

For time slivers, a weight is calculated based on the width of the timing resolution in bins of PE. In each bin of PE, the RMS of the timing residual from Figure 5.7 is computed. A parameterization to the RMS as a function of PE is determined empirically to be $1.554 + 14.458/PE$. Rather than weight by the number of photoelectrons, hits are weighted by the inverse of this fit:

$$\text{wgt} = \frac{1}{1.554 + 14.458/PE}. \quad (6.1)$$

The threshold for a time sliver is 1.5 units of this weight. A sliver is opened when the total weight in a 5 ns window exceeds 1.5. The window is slid forward in time, and closed when the threshold is no longer met. An example of an event reconstructed by finding time slivers is shown in Figure 6.7. In this event, the K^+ decays to $\pi^+\pi^0$. A kinked track is not found because the π^+ is very short. Two time slivers are identified, and indicated by hits colored blue for the earlier sliver and red for the delayed sliver. The detached red hits are due to photons from the π^0 decay.

A sliver is considered a K^+ decay product candidate if its best-fit time is at least 9 ns later than the time of the stopping K^+ track candidate. The distribution of the time gap between the sliver containing the K^+ candidate and the sliver containing the decay products is shown in Figure 6.8a. Slivers with small time gaps relative to




Figure 6.7: An event where no kinked track is reconstructed, but a K^+ is identified because the event is split into two time slivers. Hits colored in blue are in the first sliver, while red hits are in the second sliver 15 ns later. The event is from data, and the particle labels are the best guess of the author.

the track are typically due to interacting pions or protons, or due to other activity from the primary neutrino interaction. About half of true stopping kaons are rejected because the decay occurs promptly and cannot be separated from the much larger background from interacting hadrons.

6.2.3 Low-energy K^+ scan

Events can be accepted even in the absence of a K^+ track. When no stopping track is found, delayed time slivers are considered kaon decay candidates if they are at least 11 ns later than the time of the muon track, and spatially near the neutrino interaction point. This selection can be seen in Figure 6.8b. This extends the acceptance to K^+ kinetic energies below the tracking threshold, which is approximately

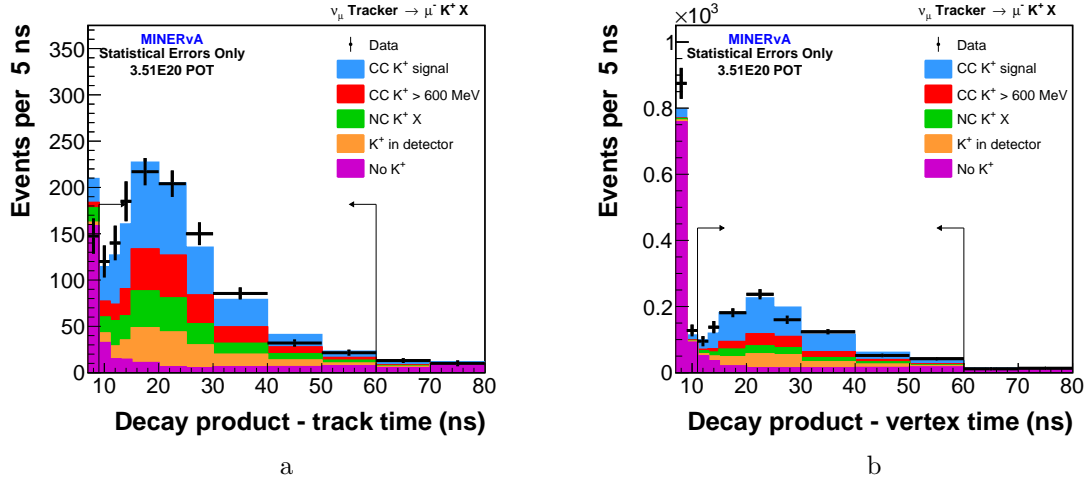


Figure 6.8: (a) The time gap between the K^+ endpoint and K^+ decay sliver for tracked events. (b) The gap between the reconstructed primary vertex time and the K^+ decay product time for vertex sliver events. In both plots, the cuts on sliver energy, number of hits, and distance are applied.

100 MeV. Events selected by this method are scanned using the Arachne event visualization program [114], and a straight line is drawn by eye connecting the start of the muon track to the nearest delayed energy deposit. The K^+ kinetic energy is estimated from the length of the line segment measured in g/cm^2 based on the simulation. An example of a “vertex sliver” event is shown in Figure 6.9.

Multiple independent visual scans were carried out on a sample of data and simulated events mixed together randomly, such that the scanner had no knowledge of whether a given event came from data or simulation. A control sample of 500 events was scanned by all scanners and used to estimate the level of disagreement between scanners in the amount of material the K^+ passed through. The average disagreement over the 500-event sample corresponds to 18 MeV of kaon kinetic energy. Potential systematic biases between individual scanners were studied and determined to be much smaller than the 18 MeV average disagreement. An uncer-




Figure 6.9: The K^+ track is not reconstructed in this event because the K^+ kinetic energy is below tracking threshold. A delayed time sliver is found spatially near the neutrino interaction point. The scanner track is shown connecting the neutrino interaction point to the nearest delayed hits (in red) that are connected by prompt hits (in blue).

tainty is added to account for differing results. The relationship between true K^+ kinetic energy and the K^+ track length as determined by the scanner in simulated events is shown in Figure 6.10. The K^+ kinetic energy residuals for tracked and scanned events are shown in Figure 6.11.

In 20% of the visually scanned events, there are no hits due to a charged particle connecting the neutrino interaction vertex and the nearest delayed hits. No kinetic energy can be estimated in these events and they are rejected. Of the rejected events in simulation, 45% are due to pile-up, in which the delayed time sliver is due to a subsequent neutrino interaction. In total, 46% of background events and only 9% of true signal events are rejected. Signal events are typically rejected when the kaon




Figure 6.10: The K^+ range as determined by scanners for simulated events, as a function of the true K^+ kinetic energy. The fitted line is used to measure the kinetic energy in data and simulation.

decay is $K^+ \rightarrow \pi^+\pi^0$ and the π^+ is obscured by prompt hits. The two π^0 photons are reconstructed, but because of the gap between the π^0 decay and photon conversion, the point where the K^+ stopped cannot be determined.

6.2.4 Additional K^+ cuts

Additional selection cuts are applied in order to reject events where the delayed energy is actually due to a “Michel” electron from the decay chain $\pi \rightarrow \mu \rightarrow e$. A K^+ at rest will decay to a μ^+ with 152 MeV of kinetic energy (and an unobserved neutrino) 64% of the time, and a back-to-back π^+ and π^0 20% of the time. Both of




Figure 6.11: Reconstructed minus true kinetic energy for tracked (left), and scanned (right) events.

these decay modes, as well as other less probable decays such as $e^+\pi^0$ will deposit approximately 150 MeV of energy in the MINERvA detector. The endpoint of the Michel electron spectrum is 55 MeV, and we select events with at least 60 MeV of reconstructed energy. The distribution of decay product energy is shown in Fig. 6.12 for events passing all other K^+ cuts. A cartoon illustrating one background rejected by this cut is shown in Figure 6.13.

In events that do not have fully-reconstructed kinked tracks, we require that the delayed time sliver have hits in at least 10 different strips. This requirement rejects events due to neutrons, which at kinetic energies of a few 100s MeV are non-relativistic, and do not lose energy by ionization as they pass through the detector. When these neutrons interact, they often produce low-energy knock-out protons late in time. These events typically produce large energy deposits in a small number of strips. The distribution of number of hits, for events passing all other K^+ selection cuts, is shown in Figure 6.14. A cartoon illustrating the background rejected by this

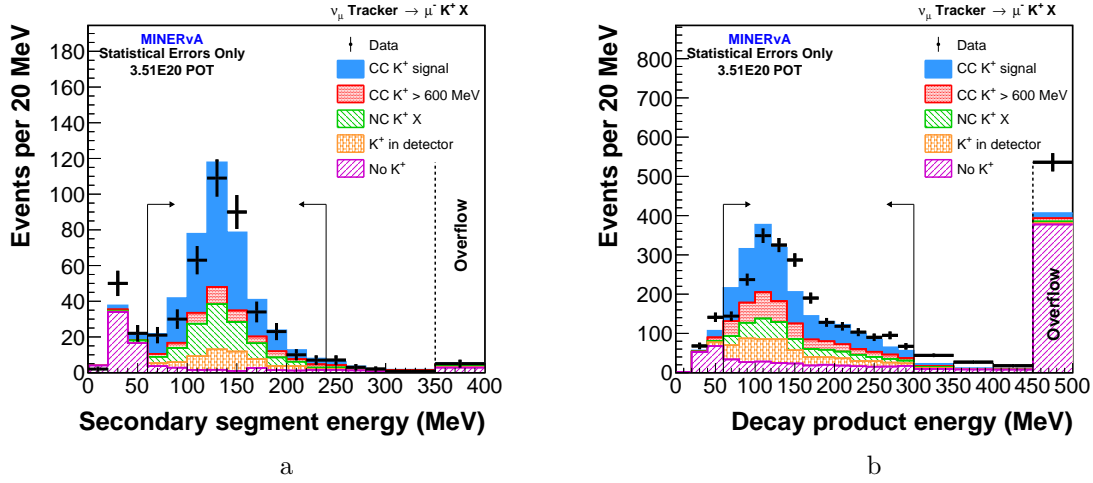


Figure 6.12: The energy of the K^+ decay product candidate when it is tracked (left) and untracked (right). Below 60 MeV, the largest background in both cases is due to Michel electrons. At high energies in the right (untracked) plot, the background is predominantly due to pile-up, which is undersimulated by 21%. The arrows show the selection criteria.

cut is shown in Figure 6.15.

Pile-up from multiple neutrino interactions in the same $10 \mu\text{s}$ beam pulse can fake the timing signature of a K^+ decay at rest. To reduce this background, we require the mean distance from the kaon endpoint vertex to a hit in the delayed time sliver to be less than 80 cm. The largest contribution to the pile-up is due to neutrino interactions in the side hadronic calorimeter that leak energy into the inner detector. Background events due to pile-up are tuned based on data in the sideband region with distance greater than 120 cm. The distribution of distance to the decay sliver is shown before and after the sideband tuning in Figure 6.16. A cartoon illustrating this background is shown in Figure 6.17.




Figure 6.13: Michel electrons from $\pi \rightarrow \mu \rightarrow e$ are typically much later than the pion track due to the $2.2 \mu\text{s}$ muon lifetime, but fast Michels have the timing signature of a K^+ decay. They are rejected due to their low energy.

6.3 Charged-current analysis

Additional cuts are made to reject neutral-current K^+ production, events where a K^+ originates from interactions inside the detector, and events where the K^+ kinetic energy is poorly reconstructed due to interactions. The flux-integrated differential cross section per nucleon in bin i is

$$\left(\frac{d\sigma}{dT_K} \right)_i = \frac{\sum_j U_{ij} (N_j - N_j^{bg})}{\epsilon_i N_{nuc} \Phi \Delta_i}, \quad (6.2)$$

where j is the index of a reconstructed T_K bin, U_{ij} is the unsmearing matrix, N_j is the number of selected events, N_j^{bg} is the predicted number of background events, ϵ_i is the selection efficiency for signal events, N_{nuc} is the number of nucleons in the




Figure 6.14: The number of hits in the delayed time sliver for events passing all other cuts, including the minimum of 60 MeV. The peak at small number of hits is primarily due to neutron interactions.

fiducial volume, Φ is the integrated ν_μ flux prediction, and Δ_i is the width of bin i .

6.3.1 CC-specific selection cuts

Charged-current events are selected by requiring that a track other than the K^+ candidate traverse more than 250 g/cm^2 of material in MINERvA, where the side and downstream calorimeters are included. Events with muons below 500 MeV of kinetic energy are rejected, 10.7% of the sample in simulation. The distribution of true muon energy for simulated events passing all other selection cuts is shown in Figure 6.18. The distribution of the longest track range is shown in Figure 6.19.




Figure 6.15: A neutron interaction that knocks out a proton can produce a small number of high-ionization hits late in time.

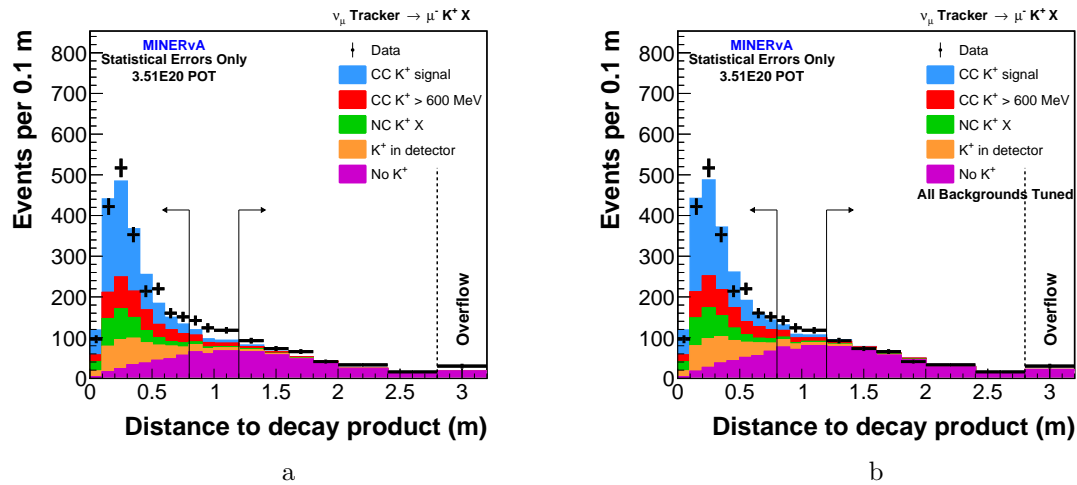


Figure 6.16: The mean distance from the K^+ track endpoint, or primary vertex in the case of untracked K^+ events, to a hit in the delayed time sliver before (left) and after (right) background tuning of the pile-up events.

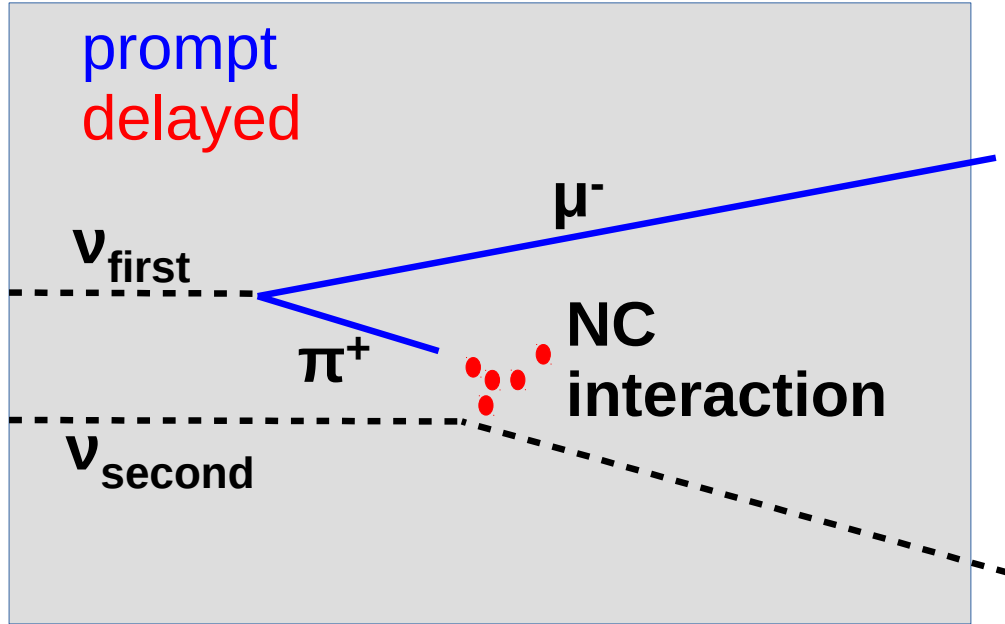


Figure 6.17: Two neutrino interactions that occur 10-60 ns apart in time can combine to look like a K^+ decay-at-rest. This background is measured in data by looking for events where the two interactions are far apart in space.

For muons that are matched with tracks in MINOS (42% of the sample), we require the curvature to be consistent with a negatively-charged particle to remove antineutrino-induced events. The charge divided by momentum Q/P is the curvature of the track in MINOS. We require that it be less than 0.1 as shown in Figure 6.20. High-momentum tracks bend less in the MINOS magnetic field, and are therefore more likely to be misreconstructed with the wrong charge. The cut is conservative, designed to retain high-momentum right-sign muons. In the simulation, 3.9% of muons that are not matched into MINOS are μ^+ from antineutrino events, and are subtracted based on the prediction from GENIE. The GENIE prediction for wrong-sign muons at lower momentum (higher Q/P) shows good agreement with data, as the data/simulation ratio in the region between 0.1 and 0.5 is 0.90 ± 0.11 .




Figure 6.18: True muon kinetic energy for CC K^+ events passing all selection cuts except the requirement that the longest track traverse more than 250 g/cm^2 of material. The first bin is events that will fail this cut. The last bin is overflow and shows the total number of muons with true kinetic energy above 15 GeV.

Hadronic interactions of high-energy charged pions can produce K^+ , for example $\pi^+ n \rightarrow K^+ \Lambda$. The K^+ can then stop and decay in the detector and mimic the signal. When this process takes place inside the nucleus of the neutrino interaction, the event is considered signal. However, when it occurs elsewhere in the detector, it must be subtracted. These events produce large hadronic showers, with an average pion energy of 3.3 GeV according to the simulation. To remove these events, we sum the hadronic energy in the detector, excluding the K^+ track. This energy includes the particles produced in the neutrino interaction, as well as products of




Figure 6.19: The range of the longest track excluding the K^+ in g/cm^2 , for events passing all other selections.

their subsequent hadronic interactions. 31% of such events are rejected by requiring that the calorimetrically-corrected hadronic energy be less than 8 GeV.

High-energy kaons which interact hadronically inside the detector are misreconstructed at much lower kinetic energy. As in the case of $\pi^+ \rightarrow K^+$ interactions, high-energy hadronic showers are produced, and 24% of interacting K^+ with true kinetic energy > 600 MeV are rejected by the cut on non- K^+ hadronic energy, which includes the products of the K^+ interaction. A sideband of events with non-kaon hadronic energy > 8 GeV is used to constrain these two classes of events simultaneously as described in Section 6.3.2. In the simulation, 6.4% of signal events with




Figure 6.20: The reconstructed charge divided by the momentum, Q/P , for MINOS-matched tracks with the cut indicated at 0.1.

$T_K < 600$ MeV are rejected by the cut on non-kaon hadronic energy.

The distribution of reconstructed non-kaon hadronic energy is shown in Fig. 6.21. After all cuts are applied, 1755 events are selected in data prior to background subtraction. A summary of event selection cuts for the charged-current sample is given in Table 6.2.

6.3.2 Background subtraction

The predicted background from simulation is scaled to agree with data in two side-band regions. The background due to beam pile-up is constrained by events where




Figure 6.21: The total non-kaon hadronic energy in the detector for events passing all other CC selections. The highest bin is overflow and is not bin-width normalized.

the mean distance from the kaon endpoint to a hit in the delayed time sliver is greater than 120 cm. In this region, 86% of the events are due to pile-up. A fit is performed to determine the scale factor for the pile-up events, with other classes of events fixed. The extracted scale factor of 1.21 is applied to backgrounds caused by pile-up.

Backgrounds from K^+ production by π^+ reactions in the detector are constrained along with kaons with true kinetic energy > 600 MeV from a sideband of events with non-kaon hadronic energy greater than 8 GeV. A single scale factor of 1.08 is determined and applied to all backgrounds, except to those due to beam pile-up.

Cut	Data	MC Total	Efficiency (%)	Purity (%)
Kink likelihood ratio	500	512	1.9	50.2
Kink secondary energy	394	424	1.8	57.7
Decay sliver time gap	35577	36590	13.4	4.9
Decay sliver energy	7503	7698	9.2	15.8
Decay sliver number of hits	3826	3561	7.5	28.1
Distance to decay sliver	2369	2372	7.1	39.6
Kaon by any method	2763	2796	8.9	42.3
Longest track range	2155	2198	8.1	48.7
Non-kaon hadronic visible energy	1837	1878	7.5	53.3
Low-energy event scan	1688	1700	7.3	56.8

Table 6.2: A summary of selected events, efficiency and purity after each cut for kaons below 500 MeV of kinetic energy. The numbers shown are cumulative. The kinked track (top section) and decay bunch (middle section) selections are combined to form the final sample (bottom section).

Signal events with true K^+ kinetic energy < 600 MeV comprise 22.5% of the sideband region. This introduces a small uncertainty due to the signal normalization into the analysis. A cross section is initially extracted leaving the normalization of these events fixed. A scale factor of 0.90 ± 0.13 is computed from the ratio of the integrated cross section in data and simulation. The analysis is repeated by applying that scale factor and its associated uncertainty to signal events in the sideband region, and results in a 3% uncertainty on the final cross section.

After subtracting backgrounds, and subtracting the estimate of events with true K^+ kinetic energy > 600 MeV, there are 885 signal events in data. Backgrounds are subtracted separately for events reconstructed by tracking and by the event scan, and the background-subtracted samples are then combined. The kinetic energy distributions for selected events with tuned backgrounds are shown for tracked and scanned events in Fig. 6.22. A breakdown of background events in the signal and high hadronic visible energy sideband regions after all tuning factors are applied is

Category	Signal region (%)	sideband region (%)
CC K^+ , $T_K < 600$ MeV	54.2	22.5
CC K^+ , $T_K > 600$ MeV	19.5	36.6
$\pi^+ \rightarrow K^+$	5.8	15.7
$K^0 \rightarrow K^+$	5.3	10.0
$\bar{\nu}_\mu$ -induced or outside F.V.	4.6	3.9
NC K^+	3.6	6.2
Pile-up	4.8	3.9
Other	2.4	1.1

Table 6.3: The breakdown of selected events for the signal and high hadronic energy sideband regions in the simulation, expressed as a percentage of the total samples prior to sideband tuning. “Other” includes events which are truly due to slow neutrons or Michel electrons.

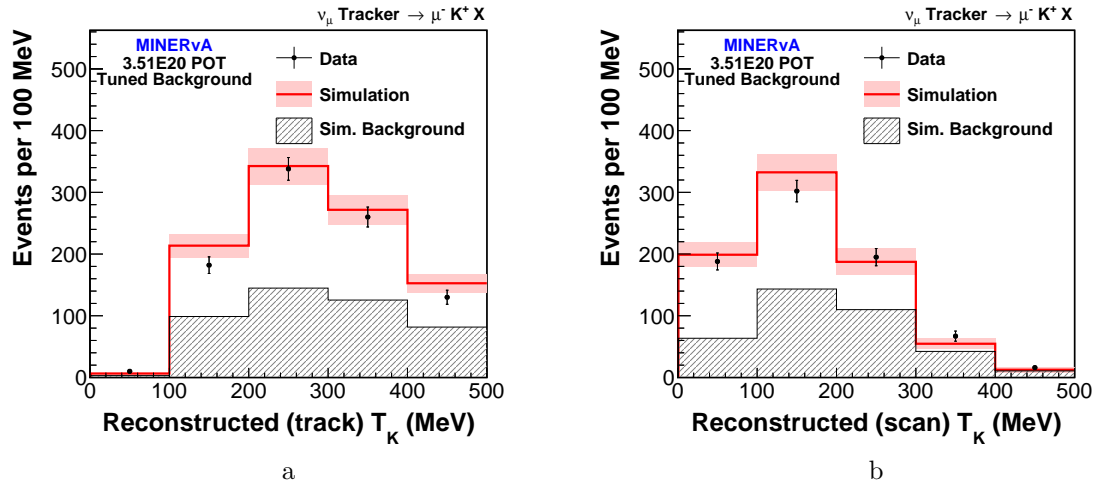


Figure 6.22: Reconstructed kinetic energy distributions with tuned background predictions for tracked (left) and untracked (right) CC events.

given in Table 6.3.

6.4 Unfolding

After background subtraction, the tracked and untracked samples are combined. The data are unfolded using a Bayesian procedure with three iterations [115]. In addition to correcting for detector resolution effects, the unfolding procedure moves events from low reconstructed kinetic energy to higher true kinetic energy because of hadronic kaon interactions in the detector. These interactions are predicted by Geant4, and reweighted to agree with external measurements of the K^+ -carbon elastic and inelastic scattering [105, 106]. The smearing matrix is shown in Fig. 6.23.

The unfolding procedure introduces correlations in the statistical uncertainties. The low and high kinetic energy regions are anticorrelated because of the feed-down from high true kinetic energy to low reconstructed kinetic energy. While we do reconstruct events in the bin from 500 to 600 MeV, 86% of the true content of that bin smears to lower reconstructed kinetic energy. This bin has a large uncertainty anticorrelated with other bins, and is not reported.

6.5 Efficiency correction

The unfolded distribution is divided by the efficiency, the integrated flux prediction, and the number of nucleons to produce the differential cross section. The efficiency for tracked and untracked events separately is shown in Figure 6.24. The tracked events include both kinked tracks and delayed time slivers near a track endpoint. The total efficiency is the sum of the two. The importance of the vertex selection in extending the acceptance to zero kinetic energy can be seen from the plots. The vertex selection has some non-zero efficiency at high kinetic energy due to kaons which interact and are reconstructed at lower kinetic energy, an effect which is




Figure 6.23: The smearing matrix, normalized so that each each row sums to unity. The number in each cell is the percentage of events of some true kaon kinetic energy that are reconstructed in a given bin.

corrected by the unfolding procedure.

A correction to the efficiency is calculated due to the lack of strangeness nonconserving reactions such as $\nu_\mu n \rightarrow \mu^- K^+ n$ in the default simulation. The reconstruction efficiency is highest for low-multiplicity final states where the delayed K^+ decay products are less likely to deposit energy in scintillator strips which have already been time-stamped with prompt energy due to other particles. Single kaon events have simple final states and thus relatively high efficiency.




Figure 6.24: Reconstruction efficiency for simulated CC K^+ events using only track selection (left) and only vertex selection (right).

The alternative simulation sample described in Section 5.1.2 is used. This sample is 8% $\Delta S = 1$ overall, but 11% $\Delta S = 1$ for $T_K < 500$ MeV. The signal efficiency in this alternate MC is higher in every bin of K^+ kinetic energy. The data are corrected by the average of the default and alternate efficiencies, with an uncertainty of 100% of the correction, such that the error band covers the difference between the efficiency obtained using the two samples. The correction is computed in each bin and is 7.0% on average. The total efficiency for the default MC is compared to the efficiency from the alternate sample in Figure 6.25.

6.5.1 Systematic uncertainties

The statistical and systematic uncertainties in each bin are given in Table 6.4. The statistical uncertainty is larger than any single systematic in every bin except for the $400 < T_K < 500$ MeV bin. The largest systematics are due to the flux, background model, and K^+ interactions in the detector.




Figure 6.25: The efficiency in the nominal simulation, along with the efficiency in the modified simulation that includes $\Delta S = 1$ events.

The uncertainty on the integrated flux is 8% [76, 77]. The sideband tuning procedure increases the uncertainty on the cross section due to the flux because the high non-kaon hadronic visible energy sideband is from the high-energy tail of the flux, while the signal region at low hadronic energy is mainly from the flux peak. While the dominant effect is from the overall flux normalization, uncertainties in the flux shape enter the analysis through the background subtraction.

The background model uncertainty is dominated by a 100% uncertainty on pion-carbon interactions inside MINERvA that produce kaons, which are simulated by

Geant4 and not constrained by external data. The sideband with high hadronic energy constrains these events, together with high-energy K^+ from signal reactions. An uncertainty arises due to the difference in the relative contribution to the signal region and non-kaon hadronic visible energy sideband region from these types of events, which can be seen in Table 6.3. Uncertainties in the GENIE FSI model are evaluated by varying its parameters within measured uncertainties [116, 117]. These variations have little effect on the analysis because the efficiency does not vary strongly with K^+ energy, and none of the significant backgrounds depend on the FSI model.

The uncertainty due to kaon interactions includes the effect on the unfolding of varying the K^+ -carbon inelastic cross section by $\pm 10\%$ to cover disagreement between the reweighted Geant4 prediction and external data shown in Figure 5.5. It also includes an uncertainty on kaon charge exchange in the detector but outside the struck nucleus. Events where K^0 production is followed by $K^0 p \rightarrow K^+ n$ are subtracted as background, while there is no acceptance for K^+ production events followed by $K^+ n \rightarrow K^0 p$. We assign a 100% uncertainty on both processes and treat it as correlated between the K^+ and K^0 charge exchange reactions.

The signal model uncertainty is due to the uncertainty in the signal rate for kaons with greater than 600 MeV of kinetic energy, and the uncertainty in the efficiency correction from single kaon production. The rate of high-energy K^+ production and the cross section for K^+ strong interactions are uncertain. We apply an uncertainty to the high-energy kaons by comparing the ratio of K^+ production cross sections above and below 600 MeV using the PYTHIA and KNO hadronization models. The nominal simulation uses KNO for $W < 2.3$ GeV, PYTHIA for $W > 3.0$ GeV, and the AGKY model in between. The resulting additional uncertainty is $^{+46}_{-11}\%$ relative

Source	0 - 100	100 - 200	200 - 300	300 - 400	400 - 500
Statistics	0.15	0.14	0.11	0.12	0.16
Flux	0.09	0.11	0.10	0.11	0.13
Background model	0.08	0.11	0.10	0.07	0.10
Kaon interactions	0.05	0.07	0.03	0.08	0.19
Signal model	0.03	0.07	0.08	0.09	0.09
Sideband tuning	0.01	0.05	0.04	0.05	0.07
Energy scale	0.02	0.02	0.03	0.04	0.05
Scanning	0.05	0.05	0.04	0.03	0.03
Total	0.21	0.24	0.21	0.22	0.33

Table 6.4: Fractional statistical and systematic uncertainties are reported in bins of kaon kinetic energy, expressed in MeV.

to the central value, where KNO, AGKY and PYTHIA are stitched together as a function of W .

The energy scale uncertainty comes from two sources. First, an uncertainty of $\pm 6\%$ is assigned to the energy of the kaon decay product to cover a discrepancy in the peak position in data relative to simulation which can be seen in Fig. 6.12. Second, uncertainties in the hadronic energy scale affect the non-kaon hadronic visible energy by pushing events from the signal to the sideband region or vice versa. We vary the detector response to hadronic, and electromagnetic, energy based on constraints from a hadron test beam, and a π^0 invariant mass peak, respectively. A summary of individual variations is given in Table 6.5. The table excludes the flux variations, which are described in Refs. [76, 77]. Also not shown are numerous variations of GENIE knob parameters, which can be found in Ref. [118].

Quantity	Variation
$\sigma(\pi^+C \rightarrow K^+X)$	$\pm 100\%$
$\sigma(K^+n \rightarrow K^0p)$	$\pm 100\%$ (correlated)
$\sigma(K^0p \rightarrow K^+n)$	
NC K^+ production rate	$\pm 30\%$
$\sigma(K^+C \rightarrow K^+X)$ inelastic	$\pm 10\%$
CC K^+ $T_K > 600$ MeV	$^{+46}_{-11}\%$
Vertex scan efficiency	$\pm 2\%$
K^+ energy response	$\pm 5\%$
Birks' constant	$\pm 13\%$
Decay product energy	$\pm 6\%$
Decay product N hits	± 1
Proton response	$\pm 4\%$
Pion response	$\pm 4\%$
Electromagnetic response	$\pm 4\%$
Neutron response	$\pm 20\%$

Table 6.5: Variations in the simulation used to estimate systematic uncertainties.

6.6 Neutral-current analysis

The neutral-current analysis follows a procedure very similar to that of the charged-current analysis presented in Section 6.3.1. This section will focus especially on the differences between the two. One important difference is that vertex sliver events are not included because there is no long muon track to use as a reliable search point for delayed time slivers. Because of this, there is no acceptance below K^+ kinetic energy of 100 MeV, and the cross section will be reported from 100 to 600 MeV of kinetic energy.

In addition to measuring $d\sigma/dT_K$, we also search explicitly for events with very little final-state energy. This is motivated by proton decay searches in water Cherenkov detectors like Super-Kamiokande, which have no sensitivity to low-energy hadrons. A cross section is reported with respect to a measure of non- K^+ energy, E_{vis} , defined as the sum of the kinetic energy of all π^\pm , K^- , and protons, and the

total energy of all photons, π^0 , and K^0 . The energy sum E_{vis} includes the prompt decay products of Λ and Σ hyperons. An attempt to extract the total hadronic energy, ν , would require using a model for the fraction of ν carried by different hadrons because the MINERvA detector responds differently to hadronic showers induced by p/π and electromagnetic showers induced by $\pi^0 \rightarrow \gamma\gamma$. Energy from neutrons is not included in the sum because neutrons are detected only when they scatter inside the detector, and most of the energy is not observed. The cross section equation is identical to 6.2, with the substitution $T_K \rightarrow E_{vis}$.

6.6.1 NC-specific selection cuts

Neutral-current events are selected by inverting the muon cut described in Section 6.3.1. Events are selected if no track, excluding the K^+ candidate, traverses more than 250 g/cm^2 of material. Figure 6.26 shows the longest track range, similar to Figure 6.19. However, in Figure 6.26, MINOS-matched tracks are not included, for reasons described in Section 6.6.2. Events with no other reconstructed tracks besides the K^+ candidate are included in the plot with a range of 0. A summary of neutral-current selection cuts is given in Table 6.6.

6.6.2 Background subtraction

The pile-up background is constrained in a manner identical to in the charged-current analysis using the sideband of time sliver events that are far from the K^+ track endpoint as shown in Figure 6.16. As in the CC analysis, backgrounds other than pile-up are constrained together. However, in the NC analysis, CC-like events are used to constrain the NC-like signal region. The sideband consists of events




Figure 6.26: The range of the longest track excluding the K^+ in g/cm^2 , for non-MINOS-matched events passing other NC selections.

with at least one track other than the K^+ , which traverses at least $250 g/cm^2$ of material in MINERvA. These events fall to the right of the cut in Figure 6.26. One important difference between the backgrounds in the two analyses is the effect of “wrong current” events. In the CC analysis, true NC K^+ events are negligible in the selected sample and do not need to be constrained. In the NC analysis, however, CC events with muon kinetic energy below 500 MeV are an important background.

True CC events in the NC-like region are constrained by extrapolating from high muon energy ($T_\mu > 500$ MeV) to low muon energy. There are two important differences between the sideband and signal region samples in this extrapolation.

Cut	Data	MC Total	Efficiency (%)	Purity (%)
Kink likelihood ratio	534	551	1.6	15.1
Kink secondary energy	423	456	1.6	17.5
Decay sliver time gap	18250	18429	5.8	1.6
Decay sliver energy	4533	4684	4.2	4.6
Decay sliver number of hits	2307	2216	3.4	7.8
Distance to decay sliver	1472	1479	3.1	10.6
Kaon by any method	1895	1935	4.7	12.2
Longest track range	492	506	4.1	41.2

Table 6.6: A summary of selected events, efficiency and purity after each cut for kaons below 600 MeV of kinetic energy. The numbers shown are cumulative. The kinked track (top section) and time sliver (middle section) selections are combined to form the final sample (bottom section). Signal events with K^+ kinetic energy above 600 MeV is considered a background for the purpose of reporting the purity. The vertex time sliver events are not included, which is why the number of decay sliver events is smaller here than in Table 6.2.

First, the mean neutrino energy is lower for the low- T_μ sample. The sideband sample includes events in high-energy tail of the flux, while the selected sample does not. This means that uncertainties which affect only the high-energy tail of the flux will lead to uncertainties on the background prediction. For example, $pC \rightarrow K^+ X$ in the target affects primarily the high-energy flux, and uncertainties on kaon production affect the sideband constraint. Second, the CC events in the NC-like region tend to be very inelastic, and have high inelasticity y . Uncertainties in the GENIE model which affect the y distribution in charged-current scattering also lead to uncertainties on the sideband constraint. These two effects can be seen in Figure 6.27, which shows a shape shape comparison in neutrino energy and y between the signal and sideband regions.

Backgrounds due to high-energy pion interactions in the detector and signal events with K^+ kinetic energy greater than 600 MeV are also constrained using the CC-like sideband and subtracted. The uncertainty on the cross section for processes




Figure 6.27: Shape comparisons between the signal and sideband regions in E_ν (left) and y (right).

like $\pi^+ n \rightarrow K^+ \Lambda$ affects π^+ produced in CC and NC reactions equally, leading to a correlation between reaction-induced backgrounds in the two samples. Similarly, the strength of K^+ -carbon inelastic processes which affect the rate of events with true $T_K > 600$ MeV that are reconstructed is also the same for CC and NC events. In addition, the shape of the K^+ energy spectrum in GENIE at very high kaon energies is correlated between CC and NC processes because the same hadronization models are used. Most high-energy K^+ come from PYTHIA ($W > 3.0$ GeV) for both CC and NC K^+ production.

The fractional breakdown of the backgrounds in the signal and sideband regions is shown in Table 6.7. In the signal region, 89% of events with $T_K > 600$ MeV are NC, while in the sideband region 81% are CC. The “Other” category includes pions and protons misidentified as kaons, as well as promptly-decaying Michel electrons with more than 60 MeV of reconstructed energy. These reconstruction failures are independent of the presence of a muon track and thus are constrained by the

Category	Signal region (%)	sideband region (%)
NC K^+ , $T_K < 600$ MeV	40.6	3.4
CC K^+ $T_K < 600$ MeV	19.4	53.2
CC or NC K^+ , $T_K > 600$ MeV	20.1	22.9
$\pi^+ \rightarrow K^+$	5.3	6.7
$K^0 \rightarrow K^+$	6.1	4.9
$\bar{\nu}_\mu$ -induced or outside F.V.	2.5	1.7
Pile-up	3.4	4.1
Other	1.9	2.6

Table 6.7: The breakdown of selected events for the signal and CC-like sideband regions, expressed as a percentage of the total samples prior to sideband tuning.

sideband.

6.6.3 Unfolding

As in the CC analysis, the data are unfolded using a Bayesian procedure with three iterations for both differential measurements. The smearing matrices are shown in Figure 6.28. The bin from 0 – 100 MeV in T_K has its upper edge right at the 100 MeV tracking threshold. It is included in the unfolding procedure but has very few events, and events in this true bin are almost always reconstructed at higher energy. The resolution in E_{vis} is limited by several factors. The amount of visible energy from a neutron is highly variable. No energy is counted in the sum, but when a neutron does interact it can convert a significant fraction of its kinetic energy to protons. Also, the visible energy due to a K^0 depends largely on the fate of the kaon in the detector. A K_S^0 decay to $\pi^+\pi^-$ deposits less visible energy than a decay to $\pi^0\pi^0$. Finally, interactions inside the detector, such as pion absorption or charge exchange, modify the visible energy.

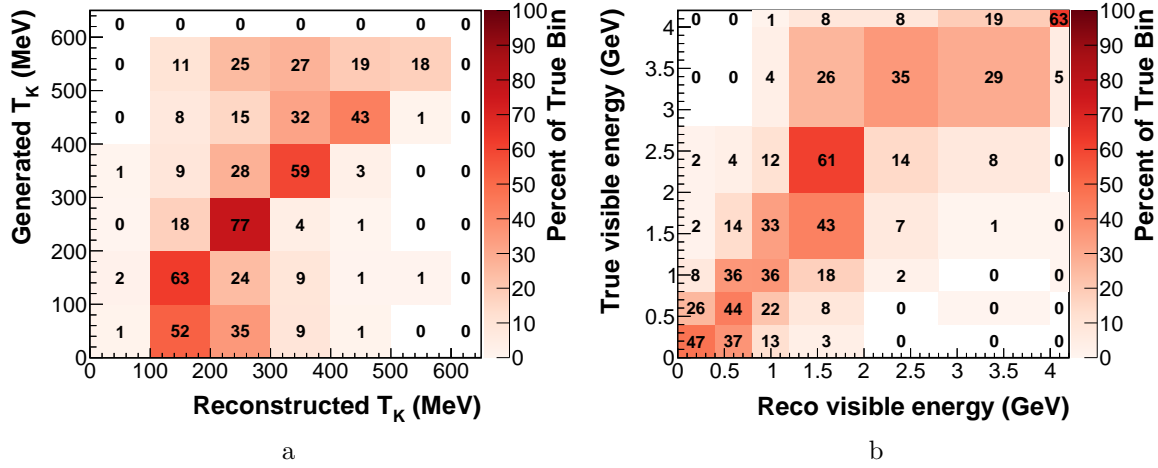


Figure 6.28: Smearing matrices for NC K^+ production in T_K (left) and E_{vis} (right).

6.6.4 Efficiency correction

The efficiency as a function of T_K and E_{vis} is shown in Figure 6.29. The efficiency with respect to T_K is for all E_{vis} . In the plot with respect to E_{vis} , T_K is cut off at 600 MeV, above which the acceptance is poor and requires the K^+ to interact inside the detector. There is no acceptance below 100 MeV of T_K due to the tracking threshold, but events below this threshold are included in the denominator. It is for this reason that the efficiency turns over at low E_{vis} – there is a correlation in the GENIE simulation between T_K and the energy of other hadrons in the final state, and events with very low total hadronic energy are less likely to have a kaon above tracking threshold.

The efficiency is low for high-multiplicity final states. In busy events, delayed energy deposits due to K^+ decay products can occur in strips where the integration window is already open by a prompt particle. In these strips, the energy of the prompt and delayed hit will be added, but the time stamp will be that of the




Figure 6.29: The NC signal efficiency as a function of T_K (left) and E_{vis} (right).

earliest photoelectron. This reduces the number of delayed hits, making it more difficult to identify the decay products of a stopped K^+ .

6.6.5 Systematic uncertainties

Summaries of fractional statistical and systematic uncertainties are given in Tables 6.8 and 6.9 for the K^+ kinetic energy and the E_{vis} measurements, respectively. The statistical uncertainty is greater than the total systematic uncertainty in every bin for both measurements. As in the CC analysis, the largest systematics are the flux, background model, kaon interactions, and the energy scale.

Again, the flux uncertainty on these cross sections is larger than the uncertainty on the integrated flux due to the background tuning procedure. Specifically, the difference in the neutrino energy spectrum for true CC events in the signal and sideband regions shown in Figure 6.27 inflates the flux uncertainty.

The background model uncertainty is dominated by the 100% uncertainty on processes like $\pi^+ C \rightarrow K^+ X$ that produce K^+ in the detector, and the uncertainty

on the CC K^+ cross section. The GENIE reweightable parameters do not significantly affect strangeness production. The parameters that specifically control final-state strange quarks described in Section 5.1.1 are not reweightable because they are used in hadronization. Therefore the default GENIE cross section uncertainty on kaon production is too small. The MINERvA CC measurement uncertainty is $\sim 20\%$ in the bins reported. It is varied by $\pm 30\%$ to cover the uncertainty on the data in addition to the overall rate disagreement with GENIE. These uncertainties are constrained by the sideband fit.

The kaon interaction uncertainty includes a 100% uncertainty K^0 and K^+ charge exchange, and a 10% uncertainty on the K^+ -carbon inelastic scattering cross section, which affects the rate of events with true energy above 600 MeV being reconstructed. These effects all affect CC events as well and are constrained by the CC-like sideband.

The uncertainty on the background model after the sideband constraint is predominantly due to the differences in the relative composition of the backgrounds in the signal and sideband regions, which is shown in Figure 6.30. Increasing the CC rate, for example, has a larger effect on the scale factor that is determined from the sideband than it does on the background prediction in the signal region.

The signal model uncertainty is due to the high-energy shape of the K^+ kinetic energy spectrum. Only kaons below 600 MeV are measured, and the uncertain rate above 600 MeV, together with the K^+ interaction cross section in the detector, contributes to the rate of high-energy kaons being misreconstructed at lower kinetic energies. This effect is small for the lowest T_K and E_{vis} bins, and large as those quantities increase. At very low T_K the contribution from kaons above 600 MeV is small, and there is a correlation between the energy of the kaon and the energy of other final-state hadrons.




Figure 6.30: Fractional breakdown of the backgrounds in the signal and sideband regions.

Source	100 - 200	200 - 300	300 - 400	400 - 500	500 - 600
Statistics	0.27	0.25	0.26	0.30	0.44
Background model	0.08	0.06	0.09	0.14	0.29
Flux	0.09	0.09	0.07	0.06	0.08
Kaon interactions	0.13	0.07	0.02	0.12	0.10
Energy scale	0.06	0.03	0.08	0.05	0.06
Sideband tuning	0.03	0.03	0.05	0.05	0.09
Signal model	0.00	0.00	0.01	0.07	0.19
Total	0.33	0.28	0.30	0.37	0.58

Table 6.8: Neutral-current systematic uncertainties on the measurement of $d\sigma/dT_K$.

The energy scale uncertainty comes from several sources. First, the energy of hits on the K^+ track is varied by $\pm 5\%$, and Birks' constant is varied by 0.015. This affects the K^+ kinetic energy reconstruction but not E_{vis} . Second, the energy of K^+ decay products is varied by $\pm 6\%$ to cover a difference between data and simulation seen in Figure 6.12. Finally, the energy responses due to protons, pions, neutrons, and electromagnetic energy are varied independently based on constraints from a miniature version of MINERvA exposed to a beam of electrons and hadrons

Source	0.0 - 0.4	0.4 - 0.8	0.8 - 1.2	1.2 - 2.0	2.0 - 2.8	2.8 - 4.0
Statistics	0.18	0.14	0.22	0.39	0.57	0.50
Background model	0.05	0.05	0.09	0.22	0.33	0.24
Flux	0.07	0.08	0.07	0.08	0.11	0.09
Kaon interactions	0.06	0.05	0.04	0.05	0.09	0.06
Energy scale	0.05	0.03	0.03	0.05	0.13	0.11
Sideband tuning	0.01	0.03	0.06	0.12	0.14	0.07
Signal model	0.02	0.00	0.04	0.12	0.17	0.12
Total	0.22	0.18	0.27	0.49	0.72	0.59

Table 6.9: Neutral-current systematic uncertainties on the measurement of $d\sigma/dE_{vis}$.

of known momentum [96]. Pion, proton, and photon energy deposits are varied by 4%, while energy observed from neutron interactions is varied by 20%. This does not affect the K^+ kinetic energy measurement, but it causes a migration in E_{vis} .

A summary of individual variations is given in Table 6.5. For the NC analysis, there is a $\pm 30\%$ variation on the CC K^+ cross section.

Chapter 7

Results

7.1 Charged-current K^+ production

The charged-current differential cross section with respect to the kaon kinetic energy is shown in Fig. 7.1, along with predictions from GENIE 2.8.4 with and without FSI, and the NuWro generator [119, 120]. Our data agree best with GENIE with FSI. The χ^2 s for 5 degrees-of-freedom for GENIE with FSI, GENIE without FSI, and NuWro are 8.1, 11.2, and 27.0, respectively. The shape-only χ^2 s for 4 degrees-of-freedom are 3.5, 7.8, and 13.1.

In GENIE, the nucleon-level cross section is tuned to inclusive K^0 and Λ production data on deuterium. Strangeness nonconserving $\Delta S = 1$ events are not simulated in this prediction, but the rate of $\Delta S = 0$ production is tuned to data that does not distinguish between the two. With the addition of a $\Delta S = 1$ component in GENIE 2.10 [108], the $\Delta S = 0$ should be reduced, as our data lie 15% below the prediction. In NuWro, kaon production is not tuned to data. Kaons are produced only in hadronization using PYTHIA for all hadronic invariant masses, and are not




Figure 7.1: The differential cross section for CC K^+ production is compared to predictions from the GENIE and NuWro event generators. The dashed red line is the GENIE prediction with the FSI model turned off.

subject to FSI.

The shape of our data agree well with the GENIE prediction with final-state interactions, despite GENIE's simplistic FSI model which lacks both kaon charge exchange and kaon production by pion reactions in the nucleus. Rescattering, which moves events to lower kinetic energies, is the only channel included in GENIE 2.8.4 for K^+ , and improves the agreement with our data significantly compared to the prediction without FSI. The addition of charge exchange would decrease the rate of final-state K^+ in an isoscalar nucleus like carbon because charged kaons outnumber neutral kaons by 50% in charged-current interactions in GENIE. Kaon production




Figure 7.2: The prediction from GiBUU for $d\sigma/dT_K$ in the MINERvA flux. Modified from an original figure in Ref. [58].

by pion reactions would enhance the cross section. The GiBUU nuclear transport model [58] predicts a dramatic increase in low-energy kaons due to FSI processes, which can be seen in Figure 7.2. Either such an enhancement is actually not very large, or the nucleon-level production of K^+ would have to be modified downward even further to compensate and still describe the data.

The cross section and absolute uncertainties are given in Table 7.1. Systematic uncertainties are strongly correlated bin to bin, especially the flux uncertainty. Small anticorrelations in the statistical uncertainties arise due to the unfolding procedure. Covariance matrices for the flux uncertainty, other systematics, and statistics are given in Appendix A in Tables A.1, A.2, and A.3, respectively. The flux uncertainty is positively correlated from bin to bin. The overall correlation of other systematics is

T_K (MeV)	$d\sigma/dT_K$	Total	Statistical	Systematic
0 - 100	0.54	0.11	0.08	0.08
100 - 200	0.52	0.12	0.07	0.10
200 - 300	0.72	0.15	0.08	0.13
300 - 400	0.74	0.16	0.09	0.14
400 - 500	0.56	0.18	0.09	0.16

Table 7.1: The differential cross section with respect to K^+ kinetic energy T_K is given in units of 10^{-39}cm^2 per nucleon per GeV, as well as the total statistical and systematic uncertainties. A breakdown of the systematic uncertainty is given in Table 6.4.

also positive, but correlations are not as strong as with the flux. The K^+ interaction uncertainty gives rise to small anticorrelations in the cross section because it modifies the probability for kaons to move from high to low energy. The unfolding procedure described in Section 6.4 also introduces small anticorrelations as events are moved between bins.

Tuning the GENIE FSI model to reproduce MINERvA data on CH does not guarantee that it would reproduce a K^+ spectrum on a heavier target like argon. The dependence of FSI on A , the number of nucleons, is a difficult problem for pion production as well. MINERvA has measured the A dependence of inclusive charged-current scattering [121] and DIS [122], and plans to study π^+ production by neutrinos on its carbon, iron, and lead targets in the upstream region of the detector. This data will be valuable in constraining how the strength of, for example, pion inelastic scattering varies with A . While MINERvA will not be able to measure K^+ production directly on different nuclear targets due to limited statistics and the difficulty of reconstructing K^+ in the nuclear targets region, this result can be combined with π^+ production on CH [113], and a future measurement of the A -dependence of π^+ production to constrain kaon FSI in other nuclei.




Figure 7.3: The differential cross section $d\sigma/dT_K$ for NC K^+ production is compared to the prediction from GENIE.

7.2 Neutral-current K^+ production

The differential cross section $d\sigma/dT_K$ for NC K^+ production is shown in Figure 7.3. The shape is observed to be relatively flat between 100 and 600 MeV of T_K , with excellent agreement between GENIE and our data, a χ^2 of just 1.8 for 5 degrees of freedom (shape-only 1.1 with 4 degrees of freedom). The cross section with absolute statistical and systematic uncertainties is given in Table 7.2. The statistical uncertainty is larger than the total systematic uncertainty in every bin.

The differential cross section $d\sigma/dE_{vis}$ is shown in Figure 7.4 compared to the prediction from GENIE. The cross section with absolute statistical and systematic

T_K (MeV)	$d\sigma/dT_K$	Total	Statistical	Systematic
100 - 200	0.22	0.07	0.06	0.04
200 - 300	0.22	0.06	0.05	0.03
300 - 400	0.18	0.05	0.05	0.02
400 - 500	0.23	0.08	0.07	0.05
500 - 600	0.21	0.12	0.09	0.08

Table 7.2: The differential cross section with respect to the K^+ kinetic energy, $d\sigma/dT_K$, is given along with absolute statistical and systematic uncertainties.

E_{vis} (GeV)	$d\sigma/dE_{vis}$	Total	Statistical	Systematic
0.0 - 0.4	0.20	0.04	0.04	0.02
0.4 - 0.8	0.22	0.04	0.03	0.02
0.8 - 1.2	0.13	0.03	0.03	0.02
1.2 - 2.0	0.12	0.06	0.05	0.04
2.0 - 2.8	0.06	0.04	0.04	0.03
2.8 - 4.0	0.19	0.11	0.09	0.06

Table 7.3: The differential cross section with respect to the non- K^+ visible energy, $d\sigma/dE_{vis}$, is given along with absolute statistical and systematic uncertainties.

uncertainties is given in Table 7.3. There is some disagreement in shape, with the data tending to favor lower visible energies compared to GENIE. The disagreement is not significant, with a χ^2 of 7.3 for 6 degrees of freedom, and 6.1 for 5 degrees of freedom shape-only, both corresponding to p -values of 0.29.

The E_{vis} measurement includes only kaons with kinetic energy below 600 MeV, as events with more energetic kaons are subtracted. The acceptance for more energetic kaons is low and strongly dependent on the model of K^+ inelastic scattering in the hydrocarbon of the tracker. Furthermore, the reconstruction of E_{vis} is always poor for high-energy kaons, which are accepted only when they interact. The definition of E_{vis} excludes the kinetic energy of the K^+ , which is underestimated when the kaon interacts, leading to an overestimation of E_{vis} .

Covariance matrices for flux, other systematics, and statistics are given for both

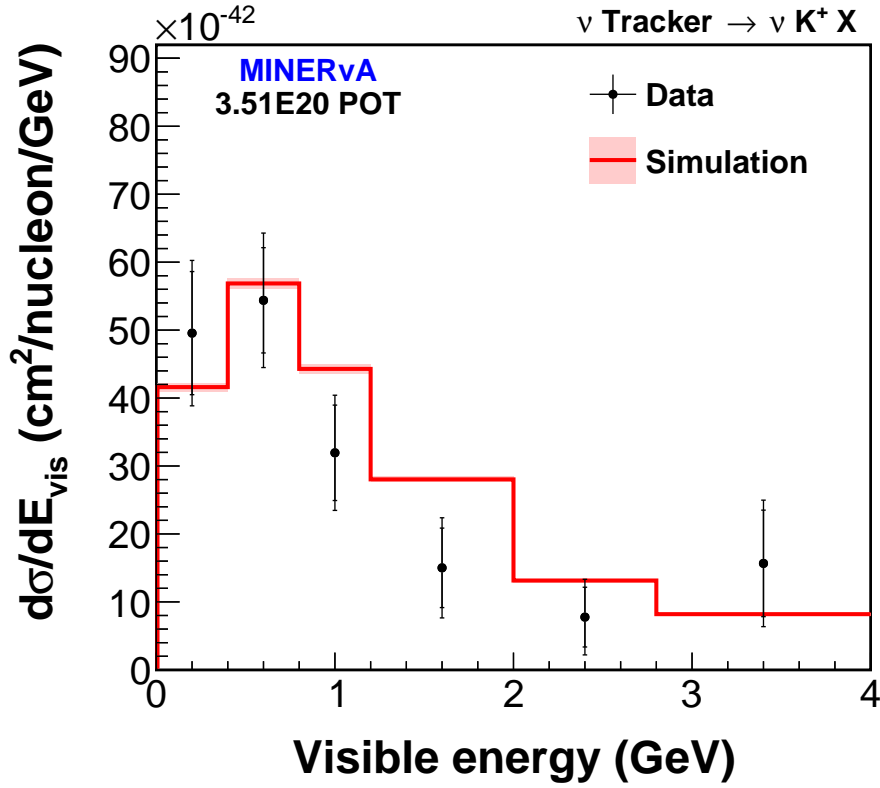


Figure 7.4: The differential cross section $d\sigma/dE_{vis}$ for NC K^+ production is compared to the prediction from GENIE.

$d\sigma/dT_K$ and $d\sigma/dE_{vis}$ in Appendix A in Tables A.4, A.5, A.6, A.7, A.8, A.9.

There are two ways in which a neutral-current neutrino-induced K^+ event would fail a $p \rightarrow K^+\nu$ selection in a water Cherenkov detector. First, additional Michel electrons would indicate the presence of final-state pions. Events with a K^+ typically have one Michel, from $K^+ \rightarrow \mu^+ \rightarrow e^+$ or $K^+ \rightarrow \pi^+ \rightarrow \mu^+ \rightarrow e^+$. Second, a photon, π^0 , or any hadron well above Cherenkov threshold would emit too much prompt light.

Figure 7.5 shows the fraction of events that are indistinguishable from a single K^+ in a water Cherenkov detector, according to GENIE. A π^0 , either in the final

state or as a decay product of a strange baryon, is assumed to always cause an event to be rejected in a proton decay search. Beyond that, four different metrics are used to determine whether other particles would cause events to be rejected. The first is to consider only events with no charged hadrons above Cherenkov threshold (see Table 2.2) and no π^+ . This assumes that the Michel efficiency is 100%, and that a hadron just above threshold will emit enough light to fail a selection on the number of prompt hits. The second metric assumes a Michel efficiency of 0%, and classifies an event as invisible if all hadrons are below threshold regardless of the presence of π^+ . The third and fourth metrics add 50 MeV/ c of momentum to the Cherenkov threshold for charged hadrons, which is approximately the ring threshold.

The fractions in Figure 7.5 are for events with reconstructed T_K below 252 MeV, the Cherenkov threshold. There is a correlation between the K^+ energy and the energy of other hadrons. Figure 7.6 shows the background-subtracted K^+ kinetic energy distribution for events with reconstructed E_{vis} less than 400 MeV, corresponding to the first bin in Figure 7.4. The spectrum is much softer than that for all E_{vis} in Figure 7.3. This is due in part to the feed-down from high true T_K to low reconstructed T_K due to K^+ interactions in the detector, as Figure 7.6 is not unfolded. Again, GENIE models the K^+ energy spectrum well, even in low-energy final states.

Figure 7.5 shows that it is rare for an invisible event to have $E_{vis} > 0.8$ GeV. We can search explicitly for π^+ and π^0 in MINERvA by looking for Michel electrons and photons, respectively. In Section 7.2.1, we look for Michel electrons in MINERvA to quantify the π^+ content of NC K^+ events at low E_{vis} . In Section 7.2.2, we look for photons to quantify the π^0 content of these events.




Figure 7.5: The fraction of events that are indistinguishable from a single K^+ to a water Cherenkov detector as a function of E_{vis} . The four different criteria for determining whether an event is single K^+ -like are described in the text.

7.2.1 π^+ content of NC K^+ events

A final-state π^+ would add an additional Michel to the event from the decay chain $\pi^+ \rightarrow \mu^+ \rightarrow e^+$. The π^+ could also originate from the decay $\Sigma^+ \rightarrow n\pi^+$. Negative hadrons typically capture on oxygen and do not produce Michels. The rate of π^+ in NC K^+ events at low E_{vis} is expected to be small due to charge conservation. Low E_{vis} events typically have small particle multiplicities, and the charge of the K^+ must be balanced, suppressing additional π^+ .

MINERvA identifies Michel electrons by searching for a 10-55 MeV energy deposit in a later time slice, which effectively requires the Michel to be at least ~ 200 ns




Figure 7.6: The reconstructed K^+ kinetic energy with backgrounds subtracted, for events with reconstructed $E_{vis} < 0.4$ GeV only.

later than the neutrino interaction. The Michel candidate must also be within 175 mm of a track endpoint, or the neutrino interaction point. Figure 7.7 shows the fraction of events with at least one reconstructed Michel in the simulation, as a function of the number of true π^+ in the final state. The prompt decay products of strange baryons, such as $\Sigma^+ \rightarrow n\pi^+$, are included. Only true NC K^+ events that pass all selection criteria are included.

Because the sample is neutral current, all Michels in Figure 7.7 come from pions, or from the K^+ decay. A 175 mm search window is centered at the endpoint of every track in the event excluding the K^+ candidate, and additionally at the neutrino




Figure 7.7: The fraction of events with at least one reconstructed Michel electron candidate, as a function of the number of true π^+ in the final state for NC K^+ events.

interaction point. The Michel is not counted when the nearest track endpoint is the K^+ or daughter μ^+ . For events without a final-state π^+ , a Michel is identified 9% of the time. This is primarily due to the Michel from the K^+ decay product, which can occur closer to another track when the K^+ interacts.

In sample with reconstructed $E_{vis} < 0.8$ GeV (the first two bins in Figure 7.4), 17% of simulated events have a final-state π^+ . This sample includes both signal and background events, with a purity of 58%. Multi- π^+ final states are negligible due to the cut on the visible energy, and the need to balance charge. Michel identification for these events is summarized in Table 7.4. The π^+ rejection efficiency is the

	All Michels	Three-view only
MC π^+ rejection efficiency	28%	12%
MC fake rate	5.5%	1.4%
MC purity	51%	64%
MC total	12.7%	3.8%
Data	$(11.8 \pm 2.6)\%$	$(4.0 \pm 1.5)\%$

Table 7.4: A summary of Michel electron identification in reconstructed NC K^+ with reconstructed $E_{vis} < 0.8$ GeV. The numbers in the table differ from Figure 7.7 because the figure is for true signal events at all E_{vis} , and the table is restricted to low E_{vis} .

fraction of events with a true final-state π^+ and a Michel electron. The fake rate is the fraction of events with no final-state π^+ and a reconstructed Michel. The purity is the fraction of events with a reconstructed Michel that also have a true π^+ . The overall rate in the simulation is compared to data.

Requiring a Michel electron to deposit energy in all three plane views reduces the accidental rate. This is because a one-view Michel only has a two-dimensional position and it is possible for activity truly distant from the search point to be near it in this two-dimensional projection. Table 7.4 shows rates for all Michels as well as only those which have hits in all three views. In both cases, the simulation is consistent with the data, indicating that the rate of final-state positive pions in NC K^+ events at low E_{vis} is well modeled. This is important for applying the MINERvA constraint to a water Cherenkov detector. In principle, the differential cross section in Figure 7.4 could be reproduced in such a way that there were many more π^+ , and fewer other particles. This would cause the background rate in a $p \rightarrow K^+\nu$ search to be higher than what Figure 7.4 would suggest. The Michel search restricts the π^+ rate in data to be $^{+45}_{-34}\%$ of the GENIE prediction at one standard deviation.




Figure 7.8: A $\nu n \rightarrow \nu K^+ \Lambda$, $\Lambda \rightarrow n \pi^0$ candidate from data, with hypotheses for the particle tracks labeled. The hits labeled as neutron are likely due to a proton ejected from the nucleus by an interacting neutron.

7.2.2 π^0 content of NC K^+ events

Figure 7.8 shows an example of a $\nu n \rightarrow \nu K^+ \Lambda$ candidate where the Λ decays to $n \pi^0$. Even at low E_{vis} , such an event would be rejected from a $p \rightarrow K^+ \nu$ search in a water Cherenkov detector due to the light produced by the photons. Figure 7.9 shows an example of an NC K^+ event with no π^0 . A visual scan was performed to identify photon candidates in data and simulation. The scan sample consisted of data and simulated events with true $E_{vis} < 0.8$ GeV, mixed randomly so that the scanner does not know whether any given event is from data or simulation.

In simulated events with true $E_{vis} < 0.8$ GeV, 35% have a true π^0 either in the final state or as a decay product of a strange baryon. In the scan, 80% of such events were marked as having an obvious detached photon, compared to just 6% of events that did not have a true π^0 . In the sample of events marked by the scanner as having a photon, 88% had a true π^0 .




Figure 7.9: An NC K^+ candidate without a π^0 from data, with hypotheses for the particle tracks labeled. The pion could be a π^- from the decay $\Sigma^- \rightarrow n\pi^-$ with an unobserved neutron.

A photon was identified in 23% of data events compared to 33% of simulated events, giving a ratio of data to simulation of $0.70^{+0.15}_{-0.13}$. There is tension between the photon content observed in data and simulation at the 2σ level. Three independent scanners arrived at the same conclusion.

In neutral-current K^+ production, the K^+ is always produced in association with another strange particle so that $\Delta S = 0$. In simulated events, the photon content is different for K^+K^- , $K^+\Lambda$, $K^+\Sigma^+$, and $K^+\Sigma^-$ events because Λ and Σ^+ decay to a π^0 and a nucleon with 0.34 and 0.52 branching fractions, respectively. In the scan, a photon was identified in 40% of $K^+\Lambda$ and $K^+\Sigma^+$ events, compared to 18% of K^+K^- and $K^+\Sigma^-$ events. This suggests that it may be possible to increase the relative fraction of kaon-kaon final states, for example, and improve agreement with MINERvA data.

7.2.3 Toward a $p \rightarrow K^+\nu$ background prediction

There are three important differences between MINERvA and a proton decay experiment for which applying this result will require model-dependent extrapolation. First, the proton decay backgrounds will not correspond exactly to T_K and E_{vis} . Reproducing Figure 7.5 with other models is an important check. Second, this result is for a CH target, while proton decay experiments will use H_2O and Ar . Other MINERvA results which compare cross sections on different nuclei provide constraints on inclusive scattering [121] and DIS [122], but not K^+ production specifically. Third, the neutrino flux from NuMI is different than the atmospheric spectrum.

The difference in the flux shape can be seen in Figure 7.10. The atmospheric flux shown is taken from Honda [24], for the Kamioka location. The flux rises rapidly at low energies of a few 100s MeV, below the threshold for associated K^+ production of 1 GeV. Multiplying the neutrino flux per m^2 by the detector mass makes it possible to directly compare the MINERvA exposure to an exposure of a larger detector to atmospheric neutrinos. The mass used for MINERvA is 5.57 tons, the tracker fiducial volume. Where the curves cross, the number of interactions of 1 GeV neutrinos in a 1 year atmospheric exposure for a 1 Megaton detector is approximately equal to the number of interactions of 1 GeV neutrinos in MINERvA. The NuMI beam is approximately 1.8×10^5 times more intense at 1 GeV, and the MINERvA detector mass is that much smaller than 1 Megaton.

The inclusive cross section is flat as a function neutrino energy in the limit where the hadronic energy goes to zero. The cross section is approximately flat with E_ν when the ratio of hadronic energy to neutrino energy, ν/E_ν is small. Hadronic energy $\nu \approx E_{vis} + T_K + m_K$, and the $p \rightarrow K^+\nu$ backgrounds come from low E_{vis} and low T_K , which is approximately $\nu \gtrsim 0.5$ GeV. We therefore expect the cross




Figure 7.10: The MINERvA flux is compared to the Honda atmospheric neutrino and antineutrino fluxes. The vertical axis is the flux multiplied by the detector mass, so that MINERvA can be compared to a Megaton-scale detector exposed to atmospheric neutrinos.

section for proton decay backgrounds to rise from threshold at 1 GeV and flatten out at ~ 5 GeV. The GENIE prediction for NC K^+ production with no π^+ or π^0 , and no hadrons with momentum more than 50 MeV/ c above Cherenkov threshold (blue points in Figure 7.5) is shown in Figure 7.11. If we assume the cross section shape in GENIE, the MINERvA exposure is approximately equivalent to ~ 35 Megaton-years. Including antineutrinos, MINERvA has about ~ 17 Megaton-years worth of low- E_{vis} NC K^+ events. While in principle NC K^+ production on argon could be constrained with atmospheric neutrinos by DUNE, the statistics would be very




Figure 7.11: The GENIE 2.8.4 prediction for the total cross section for NC K^+ production with no π^+ or π^0 , and no hadrons with momentum more than 50 MeV/ c above Cherenkov threshold. This is the final state that corresponds to the blue points in Figure 7.5.

small compared to MINERvA.

Chapter 8

Charged-current coherent K^+ production

Neutrino-induced charged-current coherent kaon production, $\nu_\mu A \rightarrow \mu^- K^+ A$, is a rare, inelastic electroweak process that brings a K^+ on shell and leaves the target nucleus intact in its ground state. This process is significantly lower in rate than neutrino-induced charged-current coherent pion production, because of Cabibbo suppression and a kinematic suppression due to the larger kaon mass. We search for such events in the scintillator tracker of MINERvA by observing the final state K^+ , μ^- and no other detector activity, and by using the kinematics of the final state particles to reconstruct the small momentum transfer to the nucleus, which is a model-independent characteristic of coherent scattering. We find the first experimental evidence for the process at 3σ significance.

8.1 Coherent K^+ theory

Charged mesons may be produced in inelastic, coherent charged-current reactions of neutrinos. This reaction is believed to occur when an off-shell W boson fluctuates to a meson. The meson is brought on the mass shell by exchange of a particle carrying no quantum numbers with the target nucleus, as illustrated in Fig. 8.1. Charged pion production through this mechanism has been observed and measured with $\mathcal{O}(10)\%$ precision on nuclei ranging from carbon to argon and across a range of neutrino (and antineutrino) energies from a few GeV to tens of GeV [123, 124, 125, 126, 127, 128, 129, 130]. If the mechanism described above is the source of these events, then there should be an analogous, Cabibbo-suppressed process in which a K^\pm meson is produced. In this chapter, the report first observation of this process is reported.

The exchange with the nucleus must leave the nucleus in its ground state for the process to be coherent. This requires that the four momentum transfer to the nucleus, $\Delta p_N \equiv p_{Af} - p_{Ai}$, satisfy the relation $|t| \equiv |(\Delta p_N)^2| \leq \hbar^2/R^2$, where R is the radius of the nucleus. Adler's theorem [131] relates the coherent meson production cross section at $Q^2 \equiv -q^2 = 0$ to the meson-nucleus elastic cross section [132, 133, 134]. In the limit of muon and meson masses $m_\mu, m_M \ll E_\nu$, where E_ν is the neutrino energy,

$$\left. \frac{d^3\sigma_{coh}}{dQ^2 dy dt} \right|_{Q^2=0} = \frac{G_F^2}{2\pi^2} f_M^2 \frac{1-y}{y} \frac{d\sigma(MA \rightarrow MA)}{d|t|}, \quad (8.1)$$

where E_M is the final state meson energy, $y = E_M/E_\nu$, and f_M is the meson decay constant. The meson-nucleus elastic cross section and its exponential decrease with increasing $|t|$ are parameterized from meson-nucleus scattering data [133, 134, 135,




Figure 8.1: Feynman diagram for coherent charged kaon production. The square of the momentum transfer to the nucleus is $|\Delta p_N|^2 = |q - p_K|^2 = |t|$.

59, 25, 136, 119]. Models must be used to extrapolate away from $Q^2 = 0$. The model most commonly used in neutrino event generators [59, 25, 136, 119] is that of Rein and Sehgal [134], which assumes no vector current and extrapolates the axial-vector current using a multiplicative dipole form factor $F_{dipole}^2(Q^2) = 1/(1 + Q^2/m_A^2)^2$ to modify Eq. 8.1. Other authors have proposed alternate extrapolations away from $Q^2 = 0$ [137, 138, 135, 139]. At low energies, modifications to Eq. 8.1 due to finite masses become important, restricting the allowed kinematics to $Q^2 \geq m_\mu^2 \frac{y}{1-y}$ and $|t| \geq \left(\frac{Q^2 + m_M^2}{2yE_\nu}\right)^2$ [140, 141]. This restriction on $|t|$ removes more phase space for kaon production than it does for pion production due to the larger meson mass, m_M . An alternate approach is to start with the Cabibbo-suppressed single kaon production cross section on nucleons [107] at low neutrino energies and calculate a coherent sum [142].

There is no model for the coherent K^+ reaction implemented in standard neutrino

interaction generators [59, 25, 136, 119]. Our ability to precisely calculate the signal reaction even in models like the ones described above is limited by the dearth of data on low energy K^+ -nucleus elastic scattering at the relevant kaon energies [143, 144]. However, a signal simulation is needed only to determine the expected energy deposited near the neutrino interaction point, and the shape of the reconstructed $|t|$ distribution. Both of these quantities are governed primarily by well-modeled detector resolutions. The overall rate of coherent K^+ production in the simulation does not affect the result.

To simulate the signal, we generate coherent pion production events with GENIE 2.8.4 [59], and recalculate hadronic kinematics while holding the leptonic kinematics and the magnitude of the four-momentum transfer to the nucleus fixed. Due to the larger kaon mass, this results in a physical solution only 26% of the time for the flux of the MINERvA experiment [4]. The cross section is also suppressed by a factor of $(f_K/f_\pi)^2 \tan^2 \theta_C = 0.077 \pm 0.001$ [5], where θ_C is the Cabibbo angle. The ratio of the K^+ -nucleus to π^+ -nucleus elastic scattering cross sections is calculated to be approximately 0.7 [145], consistent with measurements at higher meson energies [143]. This suggests that the flux-averaged cross section in our experiment for coherent K^+ production should be approximately 4×10^{-41} cm² per carbon nucleus. With 1.52×10^{29} carbon nuclei in the tracker, an integrated ν_μ flux of 2.95×10^8 neutrinos per cm² per POT, and 3.51×10^{20} POT collected, this corresponds to approximately 60 coherent K^+ events. The flux averaged total cross section for incoherent charged-current K^+ production in GENIE is 9.7×10^{40} cm² per nucleon, or about 1.2×10^{38} cm² per carbon nucleus. The cross section ratio is

$$\frac{\sigma(CC \text{ coherent } K^+)}{\sigma(CC \text{ } K^+ \text{ inclusive})} = 3.3 \times 10^{-3} \quad (8.2)$$

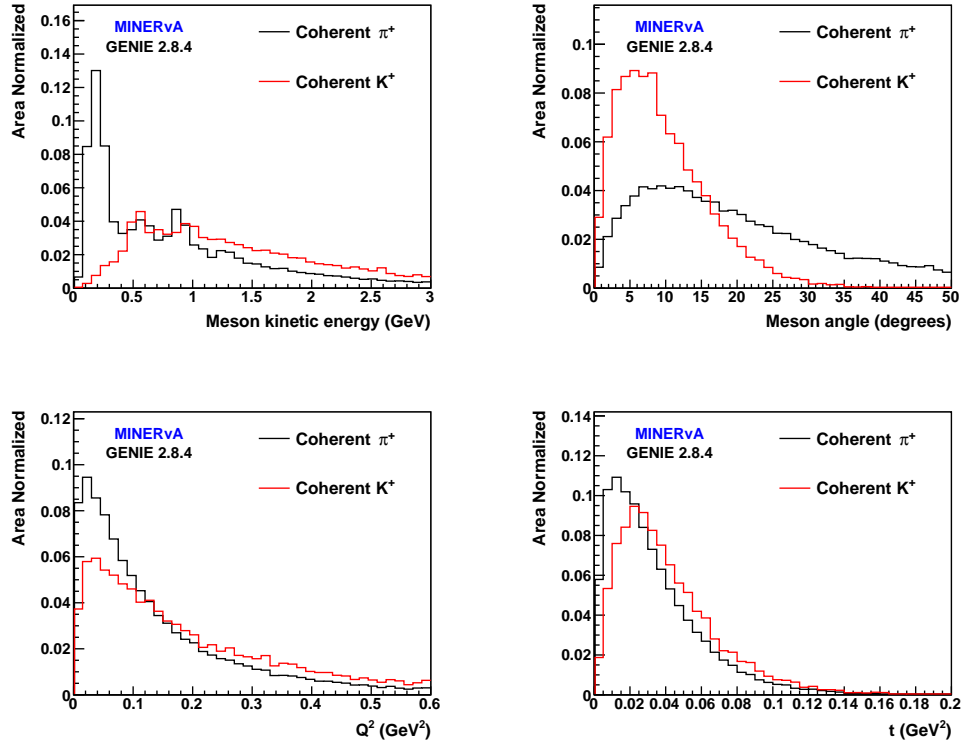


Figure 8.2: Shape comparisons for coherent pion and kaon production for meson kinetic energy, meson angle, Q^2 , and $|t|$.

in our model, meaning that about 0.3% of CC K^+ events are coherent.

Area-normalized comparisons of coherent π^+ and coherent K^+ kinematics are shown in Figure 8.2. Coherent π^+ events that give a physical solution when the meson mass is increased are preferentially higher-energy and more forward. The peak at 200 MeV of meson kinetic energy in the coherent π^+ spectrum is suppressed. The high-angle events are suppressed due to an anticorrelation between meson energy and angle. The minimum four-momentum transfer to the nucleus, t_{min} , is higher for the larger meson mass, leading to a small low- $|t|$ suppression.

8.2 Event selection

Coherent candidates must have exactly two reconstructed tracks beginning at the neutrino interaction point. One track must match to a negatively-charged track in the MINOS near detector. The other must be a K^+ candidate, passing the same K^+ selection criteria as in Section 6.2.

Events with evidence of nuclear break-up are rejected by considering energy deposited in a region around the interaction point. This “vertex region” extends 10 cm in each direction along the detector axis (3.5° above the beam direction), and 20 cm in the transverse direction. In signal reactions, energy in this region is due to the muon and kaon only. Reactions which produce additional charged hadrons deposit more energy near the vertex. Events are selected when the energy in the vertex region, E_{vtx} , is between 20 and 60 MeV, as shown in Fig. 8.4a. For the two-track sample, 75% of coherent events are retained by the vertex energy cut and 85% of two-track backgrounds are rejected. The prediction from simulation exceeds the data at very low vertex energy, consistent with previous MINERvA results [129, 146, 147], which require additional events with multiple nucleons beyond the prediction of the GENIE event generator.

The kaon energy is measured calorimetrically by summing all energy in MINERvA not associated with the muon track. For coherent events, there are no other particles in the final state so the non-muon energy is due to the K^+ and the products of its interactions in the detector. A calorimetric factor of 1.2 is applied so that the kaon energy residual in the signal simulation is peaked at zero. The reconstructed neutrino energy, E_ν , is the sum of the reconstructed muon and kaon energies, $E_\mu + E_K$. The K^+ and μ^- four-vectors are used to calculate the squared

four-momentum transferred to the nucleus:

$$\begin{aligned}
 |t| &= |\mathbf{q} - \mathbf{p}_K|^2 \\
 &= -Q^2 - 2(E_K^2 - E_\nu p_K \cos \theta_K + p_\mu p_K \cos \theta_{\mu K} + m_K^2),
 \end{aligned}
 \tag{8.3}$$

where $Q^2 = -q^2$ is the magnitude of the four-momentum transfer to the hadronic system, E_μ (E_K) is the muon (kaon) energy, p_μ (p_K) is the magnitude of the muon (kaon) three-momentum, θ_K is the kaon angle with respect to the neutrino beam, $\theta_{\mu K}$ is the angle between the outgoing kaon and muon, and m_K is the kaon mass.

The $|t|$ distributions for data and simulation are shown in Fig. 8.9. The signal simulation is normalized to the best-fit extracted from data. For simulated coherent events, 77% have reconstructed $|t| < 0.1 \text{ GeV}^2$ and 94% have reconstructed $|t| < 0.2 \text{ GeV}^2$. In simulated signal events where the kaon energy is underestimated, $|t|$ is overestimated. Typically this is due to a K^+ inelastic interaction in the detector, especially when neutrons carry a large fraction of the available energy. The shape in $|t|$ for background events is due to available phase space; according to the simulation 83% of background events have $|t| > 0.2 \text{ GeV}^2$. The $|t|$ residual is shown in Figure 8.3.

The background prediction is constrained by a sideband of events with reconstructed $|t|$ between 0.2 and 1.8 GeV^2 . In the simulation, 99.5% of the events in this region are backgrounds. Prior to the background tuning, the simulation overpredicts the background rate, with a total prediction of 23.5 events compared to 13 observed in data. We scale the simulated background by a factor of 0.55 at all $|t|$. The kaon energy distributions in simulated background events in the sideband and signal region were studied and found to be consistent.




Figure 8.3: Reconstructed - true $|t|$ for selected coherent K^+ events in the simulation. The tail where $|t|$ is overestimated is typically due to K^+ interactions in the detector where the visible energy is much less than the true K^+ energy.

Events that satisfy $20 < E_{vtx} < 60$ MeV typically have exactly two charged particles emerging from the neutrino interaction point, one μ^- and one K^+ . Background events that pass this cut may also have neutral particles that are not observed inside the vertex region. The largest single background is $\nu_\mu n \rightarrow \mu^- K^+ \Lambda$ followed by $\Lambda \rightarrow n\pi^0$. Low vertex energy events are visually scanned with Arachne [114] to remove events with an electromagnetic shower which clearly points back to the neutrino interaction point. The shower is due to the pair conversion of a photon from the decay $\pi^0 \rightarrow \gamma\gamma$. Examples of low vertex energy events with and without showers are shown in Fig. 8.5. Events are rejected only if the direction of the shower

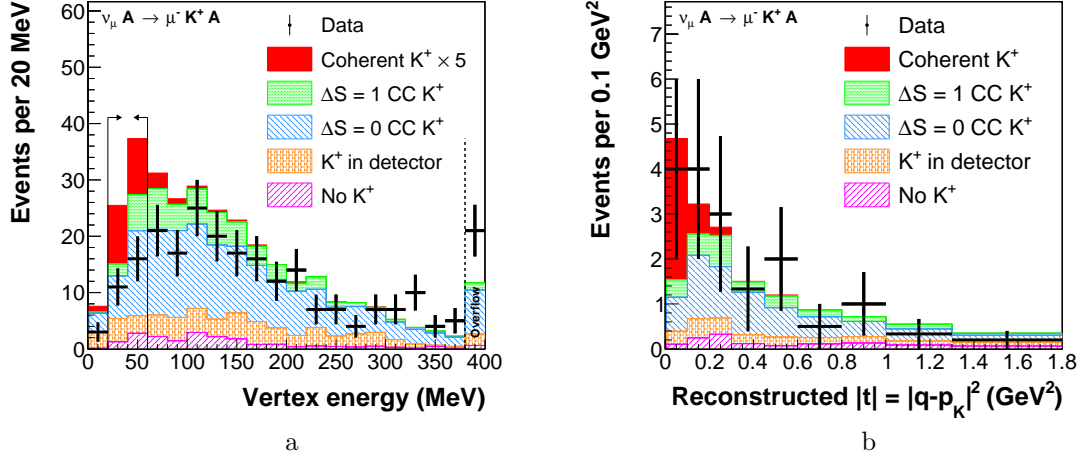




Figure 8.4: (a) The distribution of vertex energy before sideband tuning. The signal simulation is enhanced by a factor of 5. (b) The distribution of $|t|$ after selection on vertex energy and tuning of the background components, but before the removal of events with prompt π^0 s. In both figures, data points have statistical uncertainties only.

	Low E_{vtx} , $ t < 1.8$	$ t < 0.2$	No π^0 (scan)
Best-fit signal	4.05	3.80	3.77
$\Delta S = 0$	8.88 ± 2.58	2.16 ± 0.65	0.58 ± 0.20
$\Delta S = 1$ incoherent	3.34 ± 0.97	0.87 ± 0.26	0.69 ± 0.21
K^+ in detector	2.90 ± 0.84	0.78 ± 0.23	0.36 ± 0.12
No K^+	1.89 ± 0.55	0.30 ± 0.09	0.14 ± 0.05
Total sim. bkg.	17.01 ± 4.94	4.11 ± 1.20	1.77 ± 0.53
Data	21	8	6

Table 8.1: Counts for events with one μ^- and one K^+ candidate after each step in the coherent selection. The $20 < E_{vtx} < 60$ MeV cut is applied for all categories. These numbers include the scale factor of 0.55 ± 0.15 derived from the high- $|t|$ sideband. The signal is scaled to the best-fit of the unbinned likelihood fit described in the text. Backgrounds due to “ K^+ in detector” arise when π^+ or K^0 interact and produce K^+ .



a



b

Figure 8.5: (a) An event candidate in the data that is not rejected in the scan as having a π^0 candidate, compared to (b) an event that is rejected. The rejected event clearly has a photon candidate, the widest angle track in the displayed view that is detached from the neutrino interaction point and points back to the $\mu^- K^+$ vertex and not to a point of an inelastic interaction along the K^+ track. In this event display of hits in the vertically oriented scintillator strips of the detector, the neutrino beam enters from the left, and the tracks that exit out the right of the detector are μ^- .

can be determined in order to avoid removing events where the K^+ interacts inside the detector and produces a π^0 or photon.

Scanning samples contained a mixture of simulated signal and background, and

data events such that the scanner had no knowledge of whether any given event was from data or simulation. After scanning, 99% of signal events are accepted. The small inefficiency is due to π^0 produced by K^+ interactions very near the neutrino interaction point. Only 27% of $\Delta S = 0$ background events are accepted, along with 80% of $\Delta S = 1$ and 46% of other backgrounds. At low vertex energy, $\Delta S = 1$ events are dominated by $\nu_\mu n \rightarrow \mu^- K^+ n$. These events are rejected when the products of a neutron interaction in the detector point back to the neutrino interaction point. More often, the neutron interaction produces a low-energy proton for which a direction cannot be determined, and the event is accepted, leading to the relatively high acceptance for $\Delta S = 1$ backgrounds. The scan acceptance is consistent with no $|t|$ dependence. Two independent scanners agreed on signal efficiency, but differed on $\Delta S = 0$ background efficiencies. However, the differences applied to both data and simulation, and in the end, the two scan results give the same sensitivity to a coherent K^+ signal.

8.3 Results

There are roughly equal numbers of free protons and carbon nuclei in the MINERvA detector, so signal events could originate from diffractive production of K^+ from the free protons, $\nu_\mu p \rightarrow \mu^- K^+ p$, which would also produce K^+ at low $|t|$. However, many of these events would fail our selections because this reaction would produce ionization from the recoiling proton, which is much lighter than the carbon nucleus. In the limit $E_K \gg T_p$, $|t| > (\frac{m_\mu^2 \frac{E_K}{E_\mu} + m_K^2}{2E_K})^2$ and in the further limit $E_\mu \gg \frac{E_K m_\mu^2}{m_K^2}$, both of which are excellent approximations for events we accept, $|t| > \frac{m_K^4}{4E_K^2} = \frac{m_K^4}{4(m_K + T_K)^2}$. Since $T_p = |t|/2m_p$, this results in a minimum 40 MeV proton recoil energy at the

vertex for a typical T_K of 400 MeV. When combined with the K^+ and μ^- tracks this event will fail the vertex energy selection. This, plus the expected lower cross section for the production from free protons argues that low $|t|$ signal candidates are most likely coherent production from carbon. Two analyses are performed to quantify the coherent signal.

8.3.1 Cut-based analysis

In the first analysis, a cut in $|t|$ is selected based on the expected shapes for signal and background in $|t|$ to optimize sensitivity to a small signal over the predicted background. A blind study was performed to to optimize the $|t|$ cut, and the cut value was determined to be $|t| < 0.1 \text{ GeV}^2$.

To determine a p -value, pseudoexperiments are thrown. In each pseudoexperiment, a mean number of background events is calculated. Then, the number of events in that pseudoexperiment is taken from a Poisson distribution with that mean. The p -value for the null (background-only) hypothesis is the fraction of the pseudoexperiments in which the number of events is equal to or greater than the number observed in data.

The mean number of background events is obtained by throwing four random scale factors, and multiplying the product of those factors by the Monte Carlo prediction for the number of background events in the signal region. The first scale factor is the high- $|t|$ sideband constraint, equal to the number of data events in the sideband divided by the Monte Carlo prediction. The number of data events in the sideband is Poisson-fluctuated, with mean of the 13 events observed. The second scale factor is due to the statistical uncertainty on the probability to reject an event in the π^0 scan. For each of the four categories in Table 8.1, a binomial




Figure 8.6: The distributions of background scale factors, and overall acceptance of background events by the π^0 scan.

distribution is constructed based on the number of events scanned. Figure 8.6 shows the distribution of background scale factors and scan acceptance factors over many pseudoexperiments.

Two additional systematics on the background rate are considered. First, the 8% uncertainty on the integrated flux is evaluated by scaling the background by a random number in each pseudoexperiment, taken from a Gaussian with $\mu = 1.0$ and $\sigma = 0.08$. Second, an uncertainty on the background shape in $|t|$ is evaluated by considering the ratio of the Monte Carlo prediction in the sideband region to the signal region. The flux and cross section parameters are varied within their uncertainties, and the resulting spread of 6.5% in the ratio is evaluated by a second Gaussian-distributed scale factor.

The distribution of the mean number of background events in the signal region shown in Figure 8.7a. The central value background prediction is 0.69 ± 0.21 events. In each pseudoexperiment, a random number of events is taken from a Poisson distribution, where the mean is taken from the distribution in Figure 8.7a. The expected number of observed background events in many pseudoexperiments is shown




Figure 8.7: (a) The mean number of background events in many pseudoexperiments. (b) The distribution of the expected number of events in many pseudoexperiments. Three events are observed in data.

in Figure 8.7b. Three events are observed in data, corresponding to a p -value for the background-only hypothesis of 3.9%. The cut-based analysis excludes the null hypothesis at 2.1 standard deviations.

8.3.2 Likelihood analysis

The second analysis is an unbinned maximum likelihood fit [5] to the signal candidates with $|t| < 0.2 \text{ GeV}^2$. The background shape is taken from simulation, and the rate is fixed within uncertainties by the high- $|t|$ sideband constraint. The signal shape is also taken from simulation, and the only free parameter in the fit is the expected number of signal events for $|t| < 0.2 \text{ GeV}^2$. The signal and background templates used in the fit are shown in Figure 8.8.

An unbinned likelihood \mathcal{L} is defined as a function of the number of signal events μ :

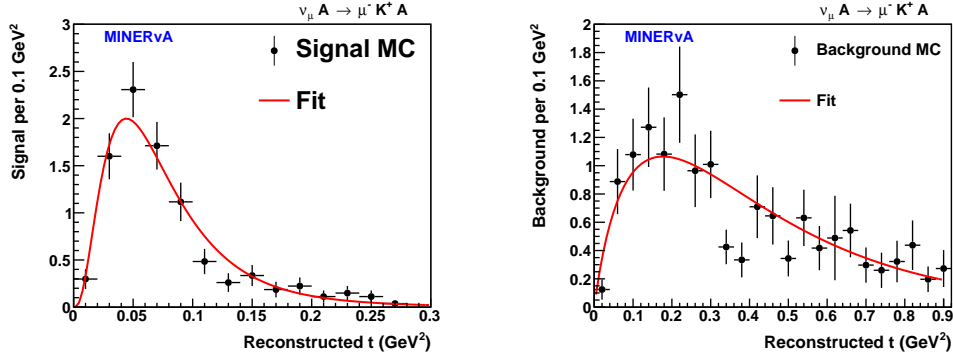


Figure 8.8: The distributions of background scale factors, and overall acceptance of background events by the π^0 scan.

$$\mathcal{L}(\mu) = \prod_{data\ i} (B(t_i) + S_\mu(t_i)) \quad (8.4)$$

where $S_\mu(t)$ and $B(t)$ are the probability density functions for signal and background, respectively, from Figure 8.8, and the signal normalization is

$$\int_0^{0.2} S_\mu(t) dt = \mu. \quad (8.5)$$

We define a ratio, R , of the likelihood of the null hypothesis ($\mu = 0$) to the likelihood at the best fit signal normalization μ_{best} :

$$R = \frac{\mathcal{L}(0)}{\mathcal{L}(\mu_{best})}. \quad (8.6)$$

Figure 8.9 shows the distribution of $|t|$ for these candidate events. The integrated number of signal events from this fit is $4.00^{+2.80}_{-2.05}$ events, where the uncertainty is the change to the number of signal events required to increase the quantity of twice the log-likelihood by one unit. The log-likelihood ratio is shown as a function of the number of coherent events in Fig. 8.9. We compare the ratio of likelihoods of the




Figure 8.9: (Left) $|t|$ for selected events after all selections, $|t| < 0.2 \text{ GeV}^2$ compared to expected distributions for the signal and background. The normalization of the background is fixed by tuning to the high- $|t|$ sideband, while the normalization of the signal is the output of the fit. The data follow the sum of the signal and background. (Right) Twice the log of the ratio of the likelihood at best fit, L_{best} , to the likelihood, L , as a function of the estimated number of coherent events. This quantity is determined by summing over the six data events and does not include systematic uncertainties.

null hypothesis of zero signal events to the best fit of 4.00, and find a log-likelihood ratio R of 7.24.

Pseudoexperiments are thrown in which the mean background rate is scaled in the same way as in the cut-based analysis described in Section 8.3.1. A random number of background events for $|t| < 0.2 \text{ GeV}^2$ is chosen from a Poisson distribution, and events are thrown based on the background template shape in Figure 8.8. The




Figure 8.10: The log-likelihood ratio, R , for pseudoexperiments with zero coherent K^+ production. The observed value in data is 7.24.

unbinned likelihood fit is performed in each pseudoexperiment, and a best-fit number of signal events is determined. The log-likelihood ratio, R , for many pseudoexperiments is shown in Figure 8.10. The value observed in data of $R = 7.24$ corresponds to a p -value of 0.3%, and a 3.0 standard deviation exclusion of the null hypothesis. The integrated number of predicted background events with $|t| < 0.2 \text{ GeV}^2$ is 1.77 ± 0.53 . Even without considering the shape, the observation of six events in data disfavors the null hypothesis.

The number of observed events predicted on the $(1.52 \pm 0.03) \times 10^{29}$ carbon nuclei in the fiducial volume of the detector can be compared with a model prediction of the cross section using information about the neutrino flux and the acceptance for




Figure 8.11: The kaon kinetic energy distribution for all accepted events in the $|t| < 0.2 \text{ GeV}^2$ signal region. Event numbers correspond to the same events in Fig. 8.9. The vertical scale on the right side corresponds to the blue curve, which represents the reconstruction efficiency as a function of true kaon kinetic energy.

coherent K^+ events, where the latter is almost completely determined by the energy of the final state kaon. The K^+ kinetic energy distribution of selected events with reconstructed $|t| < 0.2 \text{ GeV}^2$ is shown in Fig. 8.11. The six data events are more energetic than what would be expected for background events alone.

Because of finite-mass effects at low energy, $Q^2 \geq m_\mu^2 \frac{y}{1-y}$ and $|t| \geq \left(\frac{Q^2 + m_M^2}{2yE_\nu} \right)^2$ [140, 141]. The minimum $|t|$ is higher for coherent K^+ events than coherent π^+ due to the larger meson mass. The inelasticity y is E_K/E_ν in the limit of zero energy transfer to the nucleus, and the minimum $|t|$ is proportional to $(2E_K)^{-2}$. An event at low $|t|$




Figure 8.12: True $|t|$ and K^+ kinetic energy for simulated coherent signal events. The shape at low energy is due to the kinematic restriction on $|t|$.

and low E_K is therefore unlikely to be coherent. The true $|t|$ and K^+ kinetic energy in simulated signal events is shown in Figure 8.12. The K^+ energy information is not used in the extraction of the p -value of 0.3% for the background-only hypothesis.

In conclusion, coherent neutrino production of K^+ on carbon nuclei has been observed for the first time at 3σ significance by selecting events with a $\mu^- K^+$ final state, low momentum transfer to the nucleus, and no evidence of nuclear breakup. This observation, not inconsistent with our rate estimate, provides a confirmation

that the standard description of coherent meson production as arising from an off-shell W boson converting in the vacuum to a pseudoscalar meson is correct.

Chapter 9

Conclusion

The differential cross section results presented in Chapter 7 are the world's most precise measurements of charged- and neutral-current K^+ meson production by neutrinos. With 885 (201) reconstructed K^+ events after subtracting backgrounds in the charged-current (neutral-current) analysis, these results represent a significant step forward over previous measurements in bubble chambers with a few 10s of events total [66, 67, 68, 69, 70]. In addition, charged-current coherent K^+ production by neutrinos has been observed for the first time at 3.0σ significance.

These results are constraints to neutrino cross section models of K^+ production that are used to predict backgrounds in searches for the proton decay $p \rightarrow K^+\nu$. There are three important aspects of these models for proton decay searches:

1. The overall rate of NC K^+ production by neutrinos,
2. The shape of the K^+ energy distribution, and of the distribution of non-kp energy,
3. The role of K^+ -nucleus final-state interactions, which affects the K^+ energy

Cross section	χ^2 (DoF)	p-value	Shape χ^2 (DoF)	p-value	comment
CC $d\sigma/dT_K$	8.1 (5)	0.15	3.5 (4)	0.48	GENIE $\sim 15\%$ high softer E_{vis} in data
NC $d\sigma/dT_K$	1.8 (5)	0.88	1.1 (4)	0.89	
NC $d\sigma/dE_{vis}$	7.3 (6)	0.29	6.1 (5)	0.29	

Table 9.1: A summary of agreement between GENIE 2.8.4 and MINERvA data.

spectrum in $p \rightarrow K^+\nu$ events.

The neutral-current results show good agreement with GENIE in the overall rate of K^+ production. The shape of $d\sigma/dT_K$ is also in good agreement. The spectrum in E_{vis} is softer in data than in GENIE, but is statistically consistent with the prediction. These results serve as benchmarks for cross section models, which can then be used to predict backgrounds from atmospheric neutrinos in $p \rightarrow K^+\nu$ searches. A summary of the agreement between GENIE and MINERvA data is shown in Table 9.1.

The shape of $d\sigma/dT_K$ in charged-current K^+ production is sensitive to the same final-state interactions as the K^+ in a $p \rightarrow K^+\nu$ event. The agreement between GENIE and our data is improved by the addition of inelastic scattering. The overall rate is $\sim 15\%$ below the prediction, indicating that the model could be improved by either reducing the nucleon-level cross section, or perhaps by adding charge exchange, which would decrease the number of K^+ that exit the nucleus. Our data are not consistent with the prediction from NuWro, which illustrates the importance of an improved low- W DIS model for K^+ production. GENIE does not simulate K^+ in FSI by pion reactions, which adds to the cross section at low T_K in GiBUU [58]. However, no excess at low T_K is observed in data, so the nucleon-level rate must be decreased if this process is added to the FSI model to preserve agreement with MINERvA.

Bibliography

- [1] C.M. Marshall *et al.*, Phys. Rev. D **94**, 012002 (2016).
- [2] Z. Wang, C.M. Marshall *et al.*, Phys. Rev. Lett. **117**, 061802 (2016).
- [3] C.M. Marshall *et al.*, in preparation (2016).
- [4] L. Aliaga *et al.*, Nucl. Inst. Meth. A **743**, 130 (2014).
- [5] K.A. Olive *et al.* (Particle Data Group), Chin. Phys. C **38**, 090001 (2014).
- [6] H. Georgi and S.L. Glashow, Phys. Rev. Lett. **32**, 438 (1974).
- [7] J.R. Ellis, D.V. Nanopoulos, and J. Walker, Phys. Lett. B **550**, 99 (2002).
- [8] N. Arkani-Hamed, S. Dimopoulos, G.F. Giudice, and A. Romanino, Nucl. Phys. B **709**, 3 (2005).
- [9] A. Hebecker and J. March-Russell, Phys. Lett. B **539**, 119 (2002).
- [10] I. R. Klebanov and E. Witten, Nucl. Phys. B **664**, 3 (2003).
- [11] S. Dimopoulos and H. Georgi, Nucl. Phys. B **193**, 150 (1981).
- [12] N. Sakai and T. Yanagida, Nucl. Phys. B **197**, 533 (1982).
- [13] P. Nath, A. H. Chamseddine, and R. Arnowitt, Phys. Rev. D **32**, 2348 (1985).

-
- [14] Q. Shafi and Z. Tavartkiladze, Phys. Lett. B **473**, 272 (2000).
- [15] V. Lucas and S. Raby, Phys. Rev. D **55**, 6986 (1997).
- [16] J. C. Pati, Int. J. Mod. Phys. A **18**, 4135 (2003).
- [17] K. S. Babu, J. C. Pati, and F. Wilczek, Phys. Lett. B **423**, 337 (1998).
- [18] M. L. Alciati, F. Feruglio, Y. Lin, and A. Varagnolo, JHEP **03**, 054 (2005).
- [19] P. A. Cherenkov, Doklady Akademii Nauk SSSR **2**, 451 (1934).
- [20] S. Fukuda *et al.*, Nucl. Instr. Meth. A **501**, 418 (2003).
- [21] Volodymyr Takhistov (Super-Kamiokande Collaboration), arXiv:1605.03235 (2016).
- [22] K. Abe *et al.*, Phys. Rev. D **90**, 072005 (2014).
- [23] K. Abe *et al.*, Hyper-Kamokande LOI, arXiv:1109.3262 (2011).
- [24] M. Honda *et al.*, Phys. Rev. D **92**, 023004 (2015).
- [25] Y. Hayato, Acta Phys. Polon. B **40**, 2477 (2009).
- [26] R. Acciarri *et al.*, DUNE Conceptual Design Report, arXiv:1512.06148 (2015).
- [27] L. Maiani *et al.*, Phys. Rev. D **71**, 014028 (2005).
- [28] M. Ablikim *et al.*, Phys. Rev. Lett. **110**, 252001 (2013).
- [29] R. Aaij *et al.*, Phys. Rev. Lett. **115**, 072001 (2015).
- [30] M. Kobayashi and T. Maskawa, Prog. Theor. Phys. **49**, 652 (1973).

- [31] F.J. Hasert *et al.*, Phys. Lett. **46B**, 138 (1973).
- [32] S.L. Glashow, J. Iliopoulos, and L. Maiani, Phys. Rev D **2**, 1285 (1970).
- [33] Y. Fukuda *et al.*, Phys. Rev. Lett. **81**, 1562 (1998).
- [34] Q.R. Ahmad *et al.* Phys. Rev. Lett. **87**, 071301 (2001).
- [35] B. Aharmim *et al.*, Phys. Rev. C **88**, 025501 (2013).
- [36] Universitat Mainz, JUNO experiment, <http://www.staff.uni-mainz.de/wurmm/juno.html>
- [37] D.J. Gross and F. Wilczek, Phys. Rev. Lett. **30**, 1343 (1973).
- [38] H.D. Politzer, Phys. Rev. Lett. **30**, 1346 (1973).
- [39] D.I. Kazakov, arXiv:hep-ph/0012288 (2000).
- [40] S. Dimopoulos *et al.*, Phys. Lett. **112B**, 133 (1982).
- [41] J. Ellis *et al.*, Nucl. Phys. B **202**, 43 (1982).
- [42] A. Bueno *et al.*, JHEP 0704:041 (2007).
- [43] J. Hisano, H. Murayama, and T. Yanagida, Nucl. Phys. B **402**, 46 (1993).
- [44] LBNE collaboration, arXiv:1110.6249 (2010).
- [45] Duke University Saturday Academy, <http://web.archive.org/web/20110603135043/http://phy.duke.edu/saturday>
- [46] H. Nishino *et al.*, Phys. Rev. Lett. **102**, 141801 (2009).
- [47] H. Nishino *et al.*, Phys. Rev. D **85**, 112001 (2012).

-
- [48] D.S. Ayres *et al.*, NOvA Technical Design Report, FERMILAB-DESIGN-2007-01 (2007).
- [49] D. Michael *et al.*, Nucl. Instrum. Meth. A **596**, 190 (2008).
- [50] C. H. Llewellyn Smith, Phys. Rept. **3**, 261 (1972).
- [51] MC Truth, *Notorious Nucleus*, Fermilab Physics Slam (2015).
- [52] R.A. Smith and E.J. Moniz, Nucl. Phys. B **43**, 605 (1972).
- [53] A. Bodek and J.L. Ritchie, Phys. Rev. D **23**, 1070 (1981).
- [54] A. Bodek and J.L. Ritchie, Phys. Rev. D **24**, 1400 (1981).
- [55] R. Subedi *et al.*, Science **320**, 1476 (2008).
- [56] R. Gran *et al.*, Phys. Rev. D **88**, 113007 (2013).
- [57] J.T. Sobczyk, Phys. Rev. C **86**, 015504 (2012).
- [58] U. Mosel, O. Lalakulich and K. Gallmeister, Phys. Rev. D **89**, 093003 (2014).
- [59] C.Andreopoulos *et al.*, Nucl. Instr. Meth. A **614**, 87 (2010).
- [60] G.T. Jones *et al.*, Z. Phys. C **57**, 197 (1993).
- [61] D. Allasia *et al.*, Nucl. Phys. B **224**, 1 (1983)
- [62] P. Bosetti *et al.*, Nucl. Phys. B **209**, 28 (1982).
- [63] S. Willocq *et al.*, Z. Phys. C **53**, 207 (1992).
- [64] N.J. Baker *et al.*, Phys. Rev. D **34**, 1251 (1986).

-
- [65] D. DeProspero *et al.*, Phys. Rev. D **50**, 6691 (1994).
- [66] S.J. Barish *et al.*, Phys. Rev. Lett. **33**, 1446 (1974).
- [67] W. A. Mann *et al.*, Phys. Rev. D **34**, 2545 (1986).
- [68] N. J. Baker *et al.*, Phys. Rev. D **24**, 2779 (1981).
- [69] D. Son *et al.*, Phys. Rev. D **28**, 2129 (1983).
- [70] H. Deden *et al.*, Phys.Lett. B **58**, Iss. 3, 361 (1975).
- [71] P. Astier *et al.*, Nucl. Phys. B **621**, 3 (2002).
- [72] P. Adamson *et al.*, Nucl. Instr. Meth. A **806**, 279 (2016).
- [73] Pavlovich, Zarko, Ph.D. Thesis, University of Texas, FERMILAB-THESIS-2008-59 (2008).
- [74] S. Agostinelli *et al.*, Nuc. Instr. Meth. A **506**, 250 (2003).
- [75] J. Allison *et al.*, IEEE Trans. Nuc. Sci. **53** No. 1, 270 (2006).
- [76] L. Aliaga *et al.*, arXiv:1607.00704, FERMILAB-PUB-16-091-ND (2016).
- [77] Aliaga Soplin, Leonidas, Ph.D. Thesis, William & Mary, FERMILAB-THESIS-2016-03 (2016).
- [78] B. Andersson *et al.*, Nucl. Phys. B **281**, 289 (1987).
- [79] B. Nilsson-Almqvist and E. Stenlund, Comp. Phys. Comm. **43**, 387 (1987).
- [80] H. W. Bertini *et al.*, Nucl. Instr. Meth. **66**, 29 (1968).
- [81] H.W. Bertini and P. Guthrie, Nucl. Phys. A **169**, 43 (1971).

-
- [82] R.P. Feynman, Phys. Rev. Lett. **23**, 1415 (1969).
- [83] C. Alt *et al.*, Eur. Phys. J. C **49** 897 (2007).
- [84] Lebedev, Andrey V., Ph.D. Thesis, Harvard University, FERMILAB-THESIS-2007-76 (2007).
- [85] A. Ferrari, P.R.Sala, R. Guaraldi and F. Padoani, Nucl. Instr. Meth. in Phys. Res. B **71**, 412 (1992).
- [86] N. Abgrall *et al.*, Phys. Rev. C **84**, 034604 (2011).
- [87] Wolcott, Jeremy, Ph.D. Thesis, University of Rochester, FERMILAB-THESIS-2015-26 (2015).
- [88] P. Skubic, *et al.*, Phys. Rev. D **18**, 3115 (1978).
- [89] J. Paley *et al.*, Phys. Rev. D **90**, 032001 (2014).
- [90] S. M. Seun, Ph.D. thesis, FERMILAB-THESIS-2007-61 (2007).
- [91] J. Park *et al.*, Phys. Rev. D **93**, 112007 (2016).
- [92] MINERvA Technical Design Report, FERMILAB-DESIGN-2006-01 (2006).
- [93] Hamamatsu M64 PMT, <http://sales.hamamatsu.com/en/products/electron-tube-division/detectors/photomultiplier-tubes/part-r7600-00-m64.php>, 2011.
- [94] G.N. Perdue *et al.*, Nucl. Inst. Meth A **694**, 179 (2012).
- [95] Benjamin Peirce, Astronomical Journal **45**, 161 (1852).
- [96] L. Aliaga *et al.*, Nucl. Instr. Meth. A **789**, 28 (2015).

-
- [97] L. Landau, J. Phys. **8**, 201 (1944).
- [98] D. Rein and L. M. Sehgal, Annals Phys. **133**, 79 (1981).
- [99] A. Bodek, I. Park, and U.-K. Yang, Nucl. Phys. Proc. Suppl. **139**, 113 (2005).
- [100] S.A. Dytman and A.S. Meyer, AIP Conf. Proc. **1405** 213 (2011).
- [101] Z. Koba, H. B. Nielsen, and P. Olesen, Nucl. Phys. B **40**, 317 (1972).
- [102] T. Sjostrand, S. Mrenna, and P. Skands, JHEP **05**, 26 (2006).
- [103] T. Yang *et al.*, Eur. Phys. J. C **63**, 1 (2009).
- [104] A. Rubbia *et al.*, Proc. of NUINT01, KEK (Japan) (2001).
- [105] D.V. Bugg *et al.*, Phys. Rev. **168**, 1466 (1968).
- [106] E. Friedman *et al.*, Phys. Rev. C **55**, 1304 (1997).
- [107] M. Alam *et al.*, Phys. Rev. D **82**, 033001 (2010).
- [108] M. Alam *et al.*, arXiv:1512.06882 (2015).
- [109] A.B. Kaidalov, Phys. Lett. B **116**, 459 (1982).
- [110] H. Wenzel, private communication
- [111] J.B. Birks, Proc. Phys. Soc. A **64**, 874 (1951).
- [112] R. Fruhwirth, Nucl. Instr. Meth. A **262**, 444 (1987).
- [113] B. Eberly *et al.*, Phys. Rev. D **92**, 092008 (2015).
- [114] N. Tagg *et al.*, Nucl. Instr. Meth. A **676**, 44 (2012).

- [115] G. D'Agostini, Nucl. Inst. Meth. A **362**, 487 (1995).
- [116] D. Ashery *et al.*, Phys. Rev. C **23**, 2173 (1981).
- [117] T.S.H. Lee and R.P. Redwine, Ann. Rev. Nucl. Part. Sci. **52** 23 (2002).
- [118] C. Andreopolous, GENIE user manual, arXiv:1510.05494 (2010).
- [119] T. Golan, C. Juszczak and J.T. Sobczyk, Phys. Rev. C **86**, 015505 (2012).
- [120] J.A. Nowak and J.T. Sobczyk, Acta Phys. Pol. B **37**, 2371 (2006).
- [121] B.G. Tice *et al.*, Phys. Rev. Lett. **112**, 231801 (2014).
- [122] J. Mousseau *et al.*, Phys. Rev. D **93**, 071101 (2016).
- [123] H.J. Grabosch *et al.*, Z. Phys C **31**, 203 (1986).
- [124] P. Marage *et al.*, Z. Phys C **31**, 191 (1986).
- [125] P.P. Allport *et al.*, Z. Phys C **43**, 523 (1989).
- [126] P. Vilain *et al.*, Phys. Lett. B **313**, 267 (1993).
- [127] M. Hasegawa *et al.*, Phys. Rev. Lett. **95**, 252301 (2005).
- [128] K. Hiraide *et al.*, Phys. Rev. D **78**, 112004 (2008).
- [129] A. Higuera, A. Mislivec *et al.*, Phys. Rev. Lett. **113**, 261802 (2014).
- [130] R. Acciarri *et al.*, Phys. Rev. Lett. **113**, 261801 (2014).
- [131] S.L. Adler, Phys. Rev. **135**, B963 (1964).
- [132] C.A. Piketty and L. Stodolsky, Nucl. Phys. B **15**, 571 (1970).

- [133] K.S. Lackner, Nucl. Phys. B **153**, 526 (1979).
- [134] D. Rein and L.M. Sehgal, Nucl. Phys. B **223**, 29 (1983).
- [135] Ch. Berger and L.M. Sehgal, Phys. Rev. D **79**, 053003 (2009).
- [136] D. Casper, Nucl. Phys. Proc. Suppl. **112**, 161 (2002).
- [137] S.S. Gershtein, Yu.Ya. Komachenko and M.Yu Khlopov, Sov. J. Nucl. Phys. **32**, 861 (1980).
- [138] A.A. Belkov and B.Z. Kopeliovich, Sov. J. Nucl. Phys. **46**, 499 (1987).
- [139] E.A. Paschos and D. Schalla, Phys. Rev. D **80**, 033005 (2009).
- [140] D. Rein and L.M. Sehgal, Phys. Lett. B **657**, 207 (2007).
- [141] A. Higuera and E.A. Paschos, Eur. Phys. J. Plus **129**, 43 (2014).
- [142] L. Alvarez-Ruso *et al.*, Phys. Rev. C **87**, 015503 (2013).
- [143] B. Gobbi, W. Hake, J.L. Rosen and S. Shapiro, Phys. Rev. Lett. **29**, 1278 (1972).
- [144] M.M. Sternheim and R.R. Silbar, Ann. Rev. Nucl. Part. Sci. **24**, 249 (1974).
- [145] B.Z. Kopeliovich, I.K. Potashnikova, and I.I. Strakovsky, private communication
- [146] G.A. Fiorentini, D.W. Schmitz *et al.*, Phys. Rev. Lett. **111**, 022502 (2013).
- [147] P.A. Rodrigues *et al.*, Phys. Rev. Lett **116**, 071802 (2016).

Appendices

Appendix A

Covariance matrices

T_K (MeV) ($\times 10^{-80}$)	0 - 100	100 - 200	200 - 300	300 - 400	400 - 500
0 - 100	0.253	0.263	0.356	0.371	0.310
100 - 200	0.263	0.329	0.415	0.442	0.398
200 - 300	0.356	0.415	0.557	0.571	0.503
300 - 400	0.371	0.442	0.571	0.614	0.544
400 - 500	0.310	0.398	0.503	0.544	0.506

Table A.1: The covariance for the flux uncertainty for CC $d\sigma/dT_K$.

T_K (MeV) ($\times 10^{-80}$)	0 - 100	100 - 200	200 - 300	300 - 400	400 - 500
0 - 100	0.347	0.367	0.405	0.192	0.040
100 - 200	0.367	0.694	0.673	0.425	0.248
200 - 300	0.405	0.673	1.050	0.814	0.728
300 - 400	0.192	0.425	0.814	1.248	1.395
400 - 500	0.040	0.248	0.728	1.395	2.042

Table A.2: The summed covariance for all systematic uncertainties except for the flux for CC $d\sigma/dT_K$. The largest of these are due to background modelling and K^+ interactions in the detector.

T_K (MeV) ($\times 10^{80}$)	0 - 100	100 - 200	200 - 300	300 - 400	400 - 500
0 - 100	0.693	0.023	-0.072	-0.042	-0.056
100 - 200	0.023	0.498	0.021	-0.070	-0.043
200 - 300	-0.072	0.021	0.669	-0.009	-0.083
300 - 400	-0.042	-0.070	-0.009	0.801	0.066
400 - 500	-0.056	-0.043	-0.083	0.066	0.755

Table A.3: The statistical covariance for CC $d\sigma/dT_K$ is nonzero due to the unfolding procedure, which introduces small negative correlations in the statistical uncertainty from bin to bin.

T_K (MeV) ($\times 10^{80}$)	100 - 200	200 - 300	300 - 400	400 - 500	500 - 600
100 - 200	0.036	0.034	0.019	0.024	0.027
200 - 300	0.034	0.036	0.020	0.025	0.029
300 - 400	0.019	0.020	0.014	0.015	0.016
400 - 500	0.024	0.025	0.015	0.022	0.021
500 - 600	0.027	0.029	0.016	0.021	0.030

Table A.4: The flux covariance matrix for NC $d\sigma/dT_K$.

T_K (MeV) ($\times 10^{80}$)	100 - 200	200 - 300	300 - 400	400 - 500	500 - 600
100 - 200	0.121	0.059	0.022	-0.056	0.010
200 - 300	0.059	0.044	0.027	0.003	0.046
300 - 400	0.022	0.027	0.038	0.053	0.100
400 - 500	-0.056	0.003	0.053	0.210	0.300
500 - 600	0.010	0.046	0.100	0.300	0.621

Table A.5: The covariance matrix for NC $d\sigma/dT_K$ for all systematics other than flux.

T_K (MeV) ($\times 10^{80}$)	100 - 200	200 - 300	300 - 400	400 - 500	500 - 600
100 - 200	0.360	0.008	-0.022	-0.052	-0.031
200 - 300	0.008	0.285	-0.005	-0.059	-0.049
300 - 400	-0.022	-0.005	0.226	0.012	-0.037
400 - 500	-0.052	-0.059	0.012	0.473	0.045
500 - 600	-0.031	-0.049	-0.037	0.045	0.857

Table A.6: The statistical covariance matrix for NC $d\sigma/dT_K$.

E_{vis} (GeV) ($\times 10^{82}$)	0.0 - 0.4	0.4 - 0.8	0.8 - 1.2	1.2 - 2.0	2.0 - 2.8	2.8 - 4.0
0.0 - 0.4	0.122	0.141	0.070	0.029	0.020	0.028
0.4 - 0.8	0.141	0.170	0.087	0.036	0.022	0.028
0.8 - 1.2	0.070	0.087	0.051	0.023	0.013	0.014
1.2 - 2.0	0.029	0.036	0.023	0.014	0.007	0.008
2.0 - 2.8	0.020	0.022	0.013	0.007	0.007	0.008
2.8 - 4.0	0.028	0.028	0.014	0.008	0.008	0.020

Table A.7: The flux covariance matrix for NC $d\sigma/dE_{vis}$.

E_{vis} (GeV) ($\times 10^{82}$)	0.0 - 0.4	0.4 - 0.8	0.8 - 1.2	1.2 - 2.0	2.0 - 2.8	2.8 - 4.0
0.0 - 0.4	0.202	0.178	0.073	0.015	-0.011	0.018
0.4 - 0.8	0.178	0.209	0.138	0.086	0.036	0.068
0.8 - 1.2	0.073	0.138	0.174	0.160	0.101	0.128
1.2 - 2.0	0.015	0.086	0.160	0.187	0.137	0.171
2.0 - 2.8	-0.011	0.036	0.101	0.137	0.109	0.134
2.8 - 4.0	0.018	0.068	0.128	0.171	0.134	0.232

Table A.8: The covariance matrix for NC $d\sigma/dE_{vis}$ for all systematics other than the flux.

E_{vis} (GeV) ($\times 10^{82}$)	0.0 - 0.4	0.4 - 0.8	0.8 - 1.2	1.2 - 2.0	2.0 - 2.8	2.8 - 4.0
0.0 - 0.4	0.821	0.361	-0.078	-0.133	-0.069	-0.046
0.4 - 0.8	0.361	0.600	0.283	0.033	-0.032	-0.084
0.8 - 1.2	-0.078	0.283	0.493	0.252	0.073	-0.025
1.2 - 2.0	-0.133	0.033	0.252	0.342	0.196	0.081
2.0 - 2.8	-0.069	-0.032	0.073	0.196	0.193	0.142
2.8 - 4.0	-0.046	-0.084	-0.025	0.081	0.142	0.614

Table A.9: The statistical covariance matrix for NC $d\sigma/dE_{vis}$.

Appendix B

The MINERvA Collaboration

D.A. Martinez Caicedo, M.F. Carneiro, H. da Motta, A. Ghosh, K. Hurtado
Centro Brasileiro de Pesquisas Físicas

D.W. Schmitz
University of Chicago

J. Mousseau, H. Ray, D. Rimal, M. Wospakrik
University of Florida

L. Bellantoni, M. Betancourt, L. Fields, D.A. Harris, M. Kiveni, J.G. Morfín, J.
Osta, G.N. Perdue, L. Rakotondravohitra
Fermilab

A. Bravar
University of Geneva

J. Felix, A. Higuera, M.A. Ramirez, E. Valencia, G. Zavala
Universidad de Guanajuato

T. Walton
Hampton University

E. Maher

Massachusetts College of Liberal Arts

C.E. Patrick
Northwestern University

H. Schellman
Oregon State University

N. Tagg
Otterbein University

S.A. Dytman, B. Eberly, C.L. McGivern, B. Messerly, D. Naples, V. Paolone, L.
Ren
University of Pittsburgh

E. Endress, A.M. Gago
Pontificia Universidad Catolica del Peru

A. Bercellie, A. Bodek, H. Budd, T. Cai, J. Chvojka, G.A. Díaz , R. Fine, T.
Golan, J. Kleykamp, S. Manly, C.M. Marshall, K.S. McFarland, A.M. McGowan,
A. Mislivec, J. Park, P.A. Rodrigues, D. Ruterbories, J. Wolcott
University of Rochester

Nuruzzaman, R.D. Ransome, B.G. Tice
Rutgers University

O. Altinok, H. Gallagher, T. Le, W.A. Mann
Tufts University

C. Simon
University of California at Irvine

R. Gran
University of Minnesota at Duluth

C.J. Solano Salinas
Universidad Nacional de Ingeniería

W.K. Brooks, R. Galindo, J. Miller
Universidad Técnica Federico Santa María

L. Aliaga, J. Devan, M. Kordosky, J.K. Nelson, A. Norrick, D. Zhang
College of William and Mary